IBM Tivoli Monitoring
Version 6.3 Fix Pack 2

Tivoli Enterprise Portal User's Guide

IBM

IBM Tivoli Monitoring
Version 6.3 Fix Pack 2

Tivoli Enterprise Portal User's Guide

IBM

Contents

About this information

The *IBM®Tivoli® Monitoring Tivoli Enterprise Portal User's Guide* describes the Tivoli Enterprise Portal features for working with your IBM Tivoli Monitoring products.

Users of this book should be familiar with performance monitoring concepts. If you use the Tivoli Data Warehouse, you need to be familiar with the operating system that hosts the warehouse.

The document assumes no previous experience with IBM Tivoli Monitoring. To learn more about this family of products: http://www.ibm.com/software/tivoli/solutions/service-availability/index.html.

Chapter 1. Getting started

The Tivoli Enterprise Portal is the window into your Tivoli monitored environment. The portal lets you explore your enterprise in the same way that your browser lets you explore the Internet. Consult the topics here to answer the questions, *Where do I start?* and *What can I do here?*.

New in this release

Review the latest enhancements to the Tivoli Enterprise Portal.

New in Version 6.3 Fix Pack 2

The following enhancements have been made to the Tivoli Enterprise Portal for Version 6.3 Fix Pack 2.

Multiple monitor support
> The Tivoli Enterprise Portal client supports multiple monitor configurations. For example, if you have two display monitors, you can launch the Tivoli Enterprise Portal client in the secondary monitor and the dialogs, message pop-ups, and context menus are displayed in the same secondary monitor.

Creating managed system groups without assigning managed systems
> You can now create a managed system group without assigning specific managed systems. Empty managed system groups help you organize and plan for upcoming deployments before you know exactly which specific managed systems are available. You can also use empty managed system groups in test environments. For more information, see "Creating a managed system group" on page 387.

Minor changes were also made to the Tivoli Enterprise Portal for content improvements.

New features and capability enhancements were made to the Tivoli Management Services. For more information, see the following publications in the IBM Tivoli Monitoring Information Center:
> *IBM Tivoli Monitoring Administrator's Guide*
> *IBM Tivoli Monitoring Installation and Setup Guide*

New in Version 6.3

The following enhancements have been made to the Tivoli Enterprise Portal for Version 6.3.

Application Property Installation attribute group
> The Application Property Installation attribute group provides information related to the self-describing agent installation process. See "Application Property Installation attributes" on page 432.

ITM Historical Collection attribute group
> The ITM Historical Collection attribute group provides information about the active historical collections that are exporting data from a location, including the metrics about each historical collection. See "ITM Historical Collection attributes" on page 446.

Historical Export Statistics workspace

The Historical Export Statistics workspace displays the enterprise and private historical exports to a Warehouse Proxy agent instance performed by a selected managed system. This workspace is available for monitoring servers and monitoring agents. See "Historical Export Statistics workspace" on page 395.

Historical situations

Historical_Exports_Critical, Historical_Exports_Failure, Historical_Exports_Warning, and Historical_RowsCorrupted help you manage your historical exports. See "All Managed Systems situations" on page 267.

Define symbol name/value pairs for physical or logical navigator nodes

You can now define symbol name/value pairs for physical or logical navigator nodes. You can then create Take Action commands that reference these symbols. See "Navigator item properties" on page 246.

Connect points in plot charts

The plot points are connected with lines even when gaps in time are present. This feature is only available for plot charts. Use the **Always connects points** option on the **Style** tab.

Tivoli Enterprise Portal browser support for 32 bit and 64 bit browsers

The Tivoli Enterprise Portal browser client now supports both 32 bit and 64 bit Firefox and Internet Explorer browsers on both Windows and Linux platforms. Support for the most current Extended Support Release versions of the Firefox browser is now available. Support for earlier versions of Firefox prior to version 3.6 are no longer supported in the IBM Tivoli Monitoring 6.3 release.

Java™ 7 support

The Tivoli Enterprise Portal now uses IBM Java 7 as the default JRE for all deployment modes: desktop, browser, and Java WebStart. Use of IBM Java 5 is no longer supported in the IBM Tivoli Monitoring 6.3 release.

Ability to access multiple IBM Tivoli Monitoring domains from a single browser instance

It is now possible to connect the Tivoli Enterprise Portal to multiple IBM Tivoli Monitoring domains from a single browser instance. Prior to release 6.3, if you needed to manage multiple domains using the Tivoli Enterprise Portal, you were required to either launch separate browser instances (one for each domain), or use alternate client deployment modes (such as a desktop client or Java WebStart).

Minor changes were also made to the Tivoli Enterprise Portal for content improvements.

New features and capability enhancements were made to the Tivoli Management Services. For more information, see the following publications in the IBM Tivoli Monitoring Information Center:

IBM Tivoli Monitoring Administrator's Guide
IBM Tivoli Monitoring Installation and Setup Guide

Tivoli Management Services architecture

The Tivoli Enterprise Portal is based on a client-server-agent architecture.

Client

The Tivoli Enterprise Portal client is a Java-based user interface for viewing and monitoring your enterprise network. Depending on how it was installed, you can start Tivoli Enterprise Portal as a desktop application, a web application through your browser, or Java Webstart, which downloads the installable software from the Tivoli Enterprise Portal Server, installs it, and then makes it available as a desktop application.

Server

The Tivoli Enterprise Portal client connects to its application server, the *Tivoli Enterprise Portal Server*. The Tivoli Enterprise Portal Server is a collection of software services for the client that enables retrieval, manipulation and analysis of data from the monitoring agents on your enterprise. This server connects to the Tivoli Enterprise Monitoring Server, which acts as a collection and control point for alerts received from the monitoring agents, and collects performance and availability data. The main, or hub, Tivoli Enterprise Monitoring Server correlates the monitoring data collected by agents and remote servers and passes it to the Tivoli Enterprise Portal Server for presentation and your evaluation.

Agent

Tivoli Enterprise Monitoring Agents are installed on the systems whose applications or resources you want to monitor. The monitoring agent collects the monitored data, and passes it to the Tivoli Enterprise Monitoring Server to which it is connected. The client gathers the current values of the monitored properties, or attributes, and displays them in views. It can also test the values against a threshold and display an event indicator when that threshold is exceeded.

Tivoli Enterprise Portal window

Use this topic to familiarize yourself with the elements of the Tivoli Enterprise Portal window.

The Tivoli Enterprise Portal window displays information about monitored resources in your enterprise. On the left is the Navigator, which shows the arrangement of your monitored network and allows you to access information collected by different agents on your monitored systems. On the right is a workspace. The workspace can be divided into as many smaller frames, or panes, as you can reasonably fit inside the window. When you select an item in the Navigator, a new workspace opens with a set of views for that item.

The window comprises the following elements:

Title bar

In browser mode (as shown), the title bar shows the name of the workspace.

In desktop mode, the title bar shows the name of the workspace, the name and port number of the Tivoli Enterprise Portal Server, and the user name. For example, `NT Cache Details - mars:14000 - JONDO` tells us that the NT Cache Details workspace is open and the user JONDO is connected to the Tivoli Enterprise Portal Server named mars through port number 14000.

Banner

The banner is displayed when you run Tivoli Enterprise Portal in browser mode. You can replace it with your own .GIF graphic, such as your company logo. For more information, see the *IBM Tivoli Monitoring Administrator's Guide*.

Menu bar

Tivoli Enterprise Portal has a menu bar that includes the following four menus:

File   The File menu has options for working with workspaces, setting a trace, and exiting the Tivoli Enterprise Portal. The Trace Options are used only as instructed by a IBM customer support representative.

Edit   The Edit menu has editing options for workspace properties, historical data collection configuration, policies, situations, user IDs, queries, and object groups.

View   The View menu has options for opening other workspaces for a Navigator item, hiding or showing the toolbars and status bar, refreshing the data in this workspace, turning off sound for events, and for opening other Navigator views.

Help   The Help menu opens the Tivoli Enterprise Portal Help, a 10-minute tour to give you some hands-on experience, and links to the IBM website.

Also available are menus that open when you right-click an item in the Navigator or a view in the workspace.

In browser mode you can also use the browser menu bar, which is displayed just below the title bar.

If your user ID does not have View or Modify permission for a function or does not have Workspace Author Mode permission, you will not see certain items in the menus, including the pop-up menus. For example, if you have no Workspace Author Mode permission, the **Properties** menu item is not displayed.

Toolbar

Your user ID requires Workspace Author Mode permission to create and maintain workspaces, including links.

The toolbar has four tool groupings:

Returns to the previous workspace. Click the list button to see and select from a history of workspaces as they were opened. The selected workspace refreshes with the most recent sampling data from the agent. This icon is not available in browser mode. Use the **Back** tool in your browser instead.

Moves forward to the next workspace. Click the list button to see and select from a history of workspaces that were opened after this one. The selected workspace refreshes with the most recent sampling data from the agent. This icon is not available in browser mode. Use the **Forward** tool in your browser instead.

Desktop mode only: Opens a new Tivoli Enterprise Portal window. Keyboard shortcut: Ctrl + N.

If you are using the Tivoli Enterprise Portal browser client, you can also open a new window using the browser's **File > New Window** menu option. Keyboard shortcut: Ctrl + N.

Saves the current workspace properties, including any changes to views and links. Keyboard shortcut: Ctrl + S.

Opens the Properties editor for this workspace. Keyboard shortcut: Ctrl + R.

Opens the History Collection Configuration window to define and start historical collection for the specified monitoring agents and attribute groups. Keyboard shortcut: Ctrl + H.

Opens the Workflow editor to customize policies for automation. Keyboard shortcut: Ctrl + W.

Opens the Situation editor for viewing, editing, and creating situations that alert you when the conditions they describe have been met. When you use this method to open the Situation editor instead of through the Navigator item pop-up menu, the situation is not *associated* with any Navigator item. Keyboard shortcut: Ctrl + E.

Opens the Administer Users window for adding and removing user IDs and changing user permissions. Keyboard shortcut: Ctrl + U.

Opens the Query editor. When you use this method to open the Query editor (instead of opening the Query editor through the Properties editor), you can view and customize queries but you cannot assign them to table and chart views. Keyboard shortcut: Ctrl + Q.

Opens the Object Group editor for creating and working with groups, such as managed systems and situations. Keyboard shortcut: Ctrl + O.

Opens the Performance Analyzer Configuration window for creating analytical tasks. Keyboard shortcut: Ctrl + F.

Pauses or resume automatic refresh of the data in the workspace views. The tool is disabled if no refresh interval has been set for the workspace (**View → Refresh Every**). Keyboard shortcut: Shift + Esc.

Stops loading the workspace. If the workspace is set to refresh at intervals (**View → Refresh Every**), the refresh stops until the next interval. Keyboard shortcut: Shift + Esc.

Reloads the saved workspace and refreshes the data in the workspace views. Keyboard shortcut: F5.

Opens the workspace gallery panel containing selectable thumbnail images representing the workspaces that are associated with the currently selected navigator item. Images are only displayed if the workspace has been previously accessed. Keyboard shortcut: Ctrl + G.

Turns sound off or on for open events. Enabling or disabling sound for an event is controlled through the Situation editor.

Replaces the view you click next with a pie chart, the properties of which you can customize for the data you want to include.

Replaces the view you click next with a bar chart, the properties of which you can customize for the data you want to plot.

Replaces the view you next click inside with a plot chart, which you can then customize for the data you want to include.

Replaces the view you next click inside with an area chart, which you can then customize for the data you want to include.

Replaces the view you next click with a circular gauge chart, which you can then customize for the data you want to show.

Replaces the view you next click with a linear gauge chart, which you can then customize for the data you want to show.

The notepad view is a simple text editor for writing notes about the workspace.

The table view shows a column for each attribute requested by the query and one or more rows of data. You can also get table views to show data samplings over a period of time.

The message log view shows the status of events that have been opened on the entire monitored enterprise and can include up to 100 row entries at a time.

The universal message console view displays situation and policy activities as they happen, such as when a situation has been created or deleted or a policy has been activated.

The graphic view places Navigator items and their alerts as icons on a map or a picture to represent your monitored environment. Alerts show on these icons just as they do in the Navigator.

The take action view enables you to send a command to a managed system.

The terminal view starts a 3270 or 5250 session for working with z/OS® applications, or a Telnet session for working with the TCP/IP network.

The browser view opens the Tivoli Enterprise Portal integrated browser for accessing web pages and web applications.

A topology view can be added to a workspace to show the arrangement of monitored components associated with its Navigator item. Most Navigator items have at least one topology source available for the topology view, and some have multiple topology sources from which to chose.

The situation event console view shows the status of events that have been opened on this branch of the Navigator, such as the state and how long the event has been open.

The common event console view is able to integrates events from multiple event repositories: Tivoli Enterprise Monitoring Server, Tivoli Enterprise Console event server, and Netcool/OMNIbus ObjectServer.

The Tivoli Enterprise Console® event viewer integrates events from the Tivoli Enterprise Console Server with those from the Tivoli Enterprise Monitoring Server.

For performance analysis, you can specify a time span to use for the workspaces that get opened. Historical navigation mode is easily started and stopped, and you can adjust the time span.

You can hide the toolbar with **View →Toolbar**.

Each view in the workspace also has a title bar that includes these buttons:

Brings up the Properties editor for the view. This button does not display if the view has no properties.

Shows or hides the view toolbar. This button does not display if the view has no toolbar.

Splits the view horizontally into two separate views.

Splits the view vertically into two separate views.

⬜ ⬜   Maximizes the view, at which time the button becomes a **Restore** button so you can restore the view to its original size.

❌   Removes the view from the workspace.

Many views have a toolbar of functions that are applicable to that view. For example, the ⬚ **Time Span** tool is displayed for table views, bar charts, and plot charts that query attribute groups that are historical in nature or that have been configured for historical data collection.

You can turn off the display of both the view title bar and toolbar through View →View Toolbars.

Navigator

The Navigator Physical view shows the hierarchy of your monitored enterprise, from the top level (⬚ Enterprise) down to individual groupings of information collected by monitoring agents. Every item in the Navigator has at least one workspace associated with it. When you click an item in the Navigator the *default* workspace is displayed in the application window.

The Navigator provides a physical view of your monitored network, organized by operating system platform, system type, Tivoli Enterprise Monitoring product (agents), and the attribute groups from which the agents can collect information.

Small colored *event indicator* (such as ⬚ , ⚠ , and⬚) icons overlay Navigator icons when a situation (a test of certain conditions) becomes true. As you move up the Navigator hierarchy, multiple events are consolidated to show only the indicator with the highest severity.

You might see ⬚ More... indicators in the Navigator. These indicators keep the tree compact so you can see more alerts in the viewable area without having to scroll. Click ⬚ **More...** to expand that branch of the tree.

The tool for refreshing the Navigator view includes a count of the monitoring agent changes (added or removed) that are waiting to be updated. For example, six pending updates are shown as ⬚ ⁶ **Apply pending updates**.

Tivoli Enterprise Portal also comes with a Navigator Logical view that initially shows one Navigator item. To change to the Navigator Logical view, select it from the ⬚ Navigator view list ⬚ . You can also edit the Logical view and define new Navigator views for any logical hierarchy. For example, you might have a Navigator view for Manufacturing and another for Marketing.

The right border of the Navigator view has a ⬚ button that hides the Navigator and expands the adjacent workspace views to fill the gap. Click ⬚ to restore its original size. You can also click ⬚ *<Navigator name>* in the right-most section of the status bar to restore the Navigator view; or right-click it to see a list of Navigator views to select from.

Workspace

The workspace is the working area of the application window and consists of one or more views.

A *view* is a pane in the workspace, typically a chart or table showing data collected from a monitoring agent, such as the process detail running on

the UNIX operating system. A view can be split either vertically or horizontally into two separate, independent views.

Every item in the Navigator has a default workspace associated with it. Some items might have multiple workspaces, which are accessible from the 🖩 **Workspace gallery** or from the Navigator item pop-up menu.

A workspace can be linked to other workspaces. A link can be context-sensitive, whereby you click a 🖉 link in a table row or right-click a table row or a chart data series to link to more detailed information about one of the attributes in the row or data series.

Status bar

The status bar has several sections showing the following, starting from the left:

- 🕒 **Hub Time**, which is the local time at the hub monitoring server
- 🖩 **Server Available** (or offline if you see a red "X" over the server icon), which is the status of the connection to the portal server
- Workspace name, portal server name, and port number, and the user name. If the user is in workspace administration mode, ***ADMIN MODE*** is displayed.
- If the Navigator view is minimized, a 🔗 Navigator button is displayed for you to click to restore the named Navigator.

You can turn off the display of the status bar through the View menu.

Browser client differences

The Tivoli Enterprise Portal can be started as a desktop application or through your browser by entering the URL for the Tivoli Enterprise Portal Server. There are a few differences in operation between the two clients.

Tivoli Enterprise Portal can be used with either a desktop or a browser client. For a desktop client, the application software is installed on your system and you start it as you would any other application; for a browser client, the application software is installed on the web server that is integrated with the Tivoli Enterprise Portal Server and users log on from their browsers. The software is downloaded to your system the first time you log on to Tivoli Enterprise Portal, and thereafter only when software updates are available.

Portal toolbar

Tivoli Enterprise Portal functionality is the same for both types of clients except for the ⬅ **Back**, ➡ **Forward**, 🗋 **New Window**, and 🔴 **Stop** tools, which are not on the portal toolbar because they are already available browser toolbar. As well, keyboard shortcuts also used by Internet Explorer such as F1 for Help and Ctrl+H for the History Configuration window are not available in browser mode.

Workspaces have URLs

When you are logged on to the Tivoli Enterprise Portal Server through your browser, the address bar displays a URL for each workspace. You can log on and open to the workspace of your choice by adding the workspace URL to the logon URL. You can also send the URL in an email, include it in a hypertext link in an HTML page or other document, or bookmark it.

Tabbed workspaces

Every workspace has properties that control how it accessed and displayed. One of the options enables you to open the workspace in a new

window instead of replacing the previous workspace in the same window. For browsers that are enabled for tabbed web pages, the browser client uses those settings and opens the workspace in a new tab.

Launch in context and with single sign-on

You can launch products in context, such as to the IBM Tivoli Application Dependency Discovery Manager from the browser client. If your environment and user ID are configured for single sign-on, you do not need to re-authenticate as you launch into other applications.

Custom banner

The browser client has a banner image and text that you can customize for your organization. See Adding your company logo and URL in the *IBM Tivoli Monitoring Administrator's Guide* for details.

Note: As an alternative to the browser client, you can use *IBM Web Start for Java* to download the Java Web Start client from the Tivoli Enterprise Portal Server.

Monitoring agents

Tivoli Enterprise Monitoring Agents are installed where you have resources you want to monitor. The agents and the operating system, subsystem or computer that they are running on are referred to as *managed systems*.

The Navigator Physical view shows the types of agents installed and running on each managed system.

Some agents also have subagents. In such instances the agent is the *managing agent*. For example, the WebSphere® MQ Integrator Series agent has MQSI Broker subagents. Subagents are displayed in the Navigator below the managing agent and its attribute groups.

🖳 MQSI Agent - ::KQIA
- 🖥 MQSI Components
- 🖥 Product Events
- ⊙ MQSI Broker MQSIBroker1::KQIB
    - 🖥 Broker Events
    - 🖥 Flow Events
    - 🖥 Broker Information
    - 🖥 Broker Statistics
    - 🖥 Execution Group Statistics
    - 🖥 Message Flow Statistics
    - 🖥 Neighbors
    - 🖥 Topics

The agent gathers the current values of the attributes (elements) specified in a view. Gathered attribute values can also be tested against a threshold and display an event indicator when conditions exceed the threshold.

Some agents have fewer than 100 attributes, and many have several hundred. These attribute values are displayed in the table and chart views of workspaces at the system, agent, and attribute group level of the Navigator Physical view.

When an attribute value (or range of values) is specified in a situation, as a table threshold, or as a chart or table view filter, Tivoli Enterprise Portal compares the current value with the value specified and does the following:

- If the comparison in the situation is met, then the situation is true and event indicators are displayed in the Navigator.
- If the compared value exceeds the threshold specified for the table, the cell is highlighted in red, yellow, or blue.
- If the compared value meets the filter condition, it is included in the chart or table display.

Features

The Tivoli Enterprise Portal gives you a single point of control for managing the resources your applications rely on, including a range of operating systems, servers, databases, platforms, and web components.

For example, a typical IT network might have a web server on Windows, an application server and database on UNIX, and a transaction processor on CICS® on the mainframe. The Tivoli Enterprise Portal brings all these views together in a single window so that you can see when any aspect of your network is not working as expected.

Your IBM Tivoli Monitoring products use the portal interface with these major features and you can find demos for many of them on Service Management Connect on YouTube (http://www.youtube.com/ismconnect).

Feature	Description
Customizing workspaces	Tivoli Enterprise Portal presents information in a single pane of glass called a *workspace*, which consists of one or more views. Monitoring data is retrieved at regular intervals and the results sent to the workspace in the form of chart and table views. You can start monitoring activity and system status immediately with the predefined workspaces. with just a few clicks of the mouse, you can tailor your own workspaces to look at specific conditions, display critical threshold values in red, and filter incoming data so you see only what matters. See "Custom workspace options and guidelines" on page 65.

Feature	Description
Customizing workspace views	The charts display data that the monitoring agents have gathered from the systems where they are running. They can also show data from any ODBC-compliant database you write a custom query for. See "Creating a new chart" on page 104. The notepad view opens a simple text editor for writing text that can be saved with the workspace. See "Adding a notepad view" on page 168. The table view displays data that the monitoring agents have gathered from the systems where they are running. See "Adding a table view" on page 99. The message log shows the status of all situations distributed to the managed systems in your enterprise. (Your user ID requires View permission for the Event feature.) See "Adding a message log view" on page 142. The universal message console view shows situation and policy activity, and messages received as the result of universal message generation. See "Adding a Universal Message Console view" on page 145. The graphic view places Navigator items as icons on a map or picture of your choosing. (User ID requires View permission for the Event feature.) See "Adding a graphic view" on page 146. The action enables you to send a command to a managed system. (User ID requires View permission for the Action feature.) See "Creating a take action view" on page 169. The terminal view starts a 3270, 5250, or Telnet session, and enables you to write scripts for working with z/OS applications. See "Adding a terminal view to the workspace" on page 170. The browser view opens the integrated browser for accessing web pages. See "Creating a browser view" on page 166. The topology view can be added to a workspace to show the arrangement of monitored components associated with the Navigator item. See "Creating a topology view" on page 135. The situation event console view shows the status of all situations associated with items on this branch of the Navigator view, and has tools for instant filtering and event handling. (Your user ID requires View permission for the Event feature.) See "Adding a situation event console view" on page 152. The common event console view enables you to see events from the Tivoli Enterprise Monitoring Server in the same way as the situation event console, plus it incorporates events from IBM Tivoli Enterprise Console and IBM Tivoli Netcool/OMNIbus if your managed environment is configured for those servers. See "Adding a common event console view" on page 158. The IBM Tivoli Enterprise Console view displays events from the event server and can integrate them with situation events from the Tivoli Enterprise Monitoring Server. See "Adding a Tivoli Enterprise Console event viewer" on page 164.
Links between workspaces	If you often go from one workspace to another, and from one Navigator view to another, you can build a link from one to the other to get there more quickly. You can also build links that originate from a query-based view or a graphic view and use relevant data from the source to determine the link target. See "Defining links between workspaces" on page 75.

Feature	Description
Configuring historical data collection	Configure IBM Tivoli Monitoring to retain data samplings in history files and, if you have the Tivoli Data Warehouse configured, roll off to a historical database at selected intervals, from every 15 minutes to once a day. In chart and table views you can then look over what happened hours or days ago, or as far back as you have data stored to find a pattern, and you can filter out unwanted data to speed up data retrieval and to show specific areas of interest. Pruning and summarization features enable you to keep the data stored on the warehouse at a manageable level. See "Historical collection configuration" on page 221.
Creating custom queries	Every monitoring agent comes with a set of predefined queries for every attribute group. These queries tell the monitoring server what monitoring data to retrieve from the monitoring agent to display in a chart or table view. You can create your own queries to specify exactly which attributes to retrieve to the table or chart. This saves valuable resources by retrieving only the data you want and nothing more. For example, you can use filters in the query to retrieve only the records with a file size that is greater than five megabytes.

Additionally, you can write SQL queries to ODBC data sources and apply the results to any chart or table. This enables you to show Tivoli data and data from other sources (such as third-party databases) in a single workspace. See Chapter 7, "Custom queries," on page 199. |
Take Action	The take action feature lets you enter a command or select from one of the predefined commands for your product and run it on any system in your managed network either manually or automatically in response to reported conditions. See "Take action commands" on page 357.
Launch Application	The launch application feature enables you to write definitions that start applications on your computer. A launch definition can include arguments that are passed to the application when it is started. You can invoke a definition at any time from the portal interface to launch the specified application. See "Launch application" on page 361.
Condition-based alerts	You can define conditions that are tested against monitored data at regular intervals. When matching values (or values that fall within a specified range) are found, a situation event is open and a visual indicator alerts you. You can automate a response to the event. See Chapter 10, "Situations for event monitoring," on page 253.
Object groups and managed system groups	The object group editor enables you to create named lists of managed systems, historical data collection distribution lists, and situation lists.

You can create and maintain named lists of managed systems that can be applied to situation distribution lists, queries, custom Navigator managed system assignments, policies correlated by business application group. See "Grouping objects" on page 385. |
| User administration | The user administration feature is for adding and maintaining user IDs and user groups, assigning permissions for the major features, and restricting access to specific managed systems and Navigator views. See "Administer Users" on page 374. |
| Customizing the Navigator | The portal interface comes with the Navigator Physical view, which organizes your managed enterprise by operating platform and agent type, and the Navigator Logical view, which initially shows one Navigator item. You can also define Navigator views for any logical grouping, such as a departmental or site hierarchy. See "Edit Navigator view" on page 245. |

Feature	Description
Policies	The Workflow editor enables you to design sets of automated system processes, called policies, to resolve system problems. A policy performs actions, schedules work to be performed by users, or automates manual tasks. See "Automation with policies" on page 327.
Agent Builder support	The Agent Builder is a wizard that makes it easier to build custom monitoring solutions. Using mostly point and click capabilities, the Agent Builder allows you to create a monitoring agent using multiple data providers such as WMI, Perfmon, Log scraping, scripts, and process monitoring. Over time additional data providers will be added to the Agent Builder. The Agent Builder can be configured to monitor multiple remote systems, but you must create one agent instance for each remotely monitored server.
Multiple applications in one workspace	In one workspace you can build a table or chart with data from one type of monitoring agent, and another table or chart with data from a different agent, showing views from as many different agent types as are included on that branch of the Navigator.
Link between applications	You can define a link from a workspace associated with one monitoring agent to the workspace of another.
Linking between Navigators	You can link from a workspace in one Navigator view to a workspace in a different Navigator view.

Predefined workspaces, situations, and more

Every agent is designed to monitor specific *attributes* of an operating system or other type of software. With these attributes, come definitions for use with the Tivoli Enterprise Portal functions.

Use the definitions that have been designed for your product to begin monitoring and visualizing data immediately. Some definitions are ready to use; others are dormant until activated:

Workspaces
The workspaces that open when you click a Navigator item or select from the ▦ Workspace Gallery, are predefined. They provide visual feedback of real-time values from managed systems, and historical values when historical data collection has been configured. They provide a starting point for designing your own workspaces.

Queries
The predefined workspaces are populated with data gathered as a result of queries, which are also predefined. Creating your own queries from these predefined queries enables you to add or remove attributes, apply a sort order, and pre-filter the data to keep data retrieval at a manageable level and to enable you to more easily see data of interest.

Take action commands
Some agents have predefined take action commands, such as the Windows OS agent Start Service and Stop Service commands. They are available for selection and they provide examples of the kinds of actions that you might want to define for maintaining managed systems.

Situations
The tests for conditions that you want to be alerted for are available in the

predefined situations. To avoid an overwhelming number of events opening in a newly managed environment and, conversely to avoid missing important events because the comparison criteria was set too low, many of the predefined situations for an agent are not set to start automatically. They remain stopped until you either start them manually or set them to ☑ **Run at startup**. A good way to find out which situations are set to run at startup and which ones are not, is to filter the Situation editor tree with 🐾 Show Situations, which is available when the Situation editor is opened from the Navigator pop-up menu. As agents report monitoring data, and especially when that data is being collected and stored in a Tivoli Data Warehouse, you can better determine the values and circumstances that are worthy of an alert, such as high CPU activity on a transaction server at peak times.

Policies

The predefined policies that are available with some products are not set to start automatically. This is primarily because policies automate activities, as described in the policy workflow.

Managed system groups

The Tivoli Enterprise Monitoring Server and every IBM Tivoli Monitoring product has at least one predefined managed system group, indicated by an asterisk at the beginning of the list name, such as *NT_SYSTEM for the Windows OS agent. When you assign one of these managed system groups to a situation, policy, historical collection configuration, or custom Navigator, all managed systems with that agent installed are selected.

Navigator views

Some products have custom Navigator views for access to special features.

Related concepts:

Using workspaces

Custom queries
In views that display monitored data, attribute values from managed systems are retrieved by queries to the Tivoli Enterprise Monitoring Server. You can edit the queries that are used in the predefined workspaces for your monitoring products, or create new queries to populate new views. In addition, you can retrieve data from any JDBC- or ODBC-compliant database to display in a chart or table by writing an SQL SELECT statement.

Take action commands
The take action feature enables you to invoke a predefined command or enter a command on a selected managed system. Action commands can also be added to situations to implement simple (reflex) automation when an event is opened.

Situations for event monitoring
Situations are definitions of conditions to test, such as slow transaction rates or invalid logon attempts. Use situations to raise alerts of certain conditions and to trigger single action commands or automated workflows.

Automation with policies
A *policy* is a logical expression that describes a series, or *workflow*, of automated steps, called *activities*, whose order of execution you control. You can create policies that implement fully automated workflow strategies and policies that require user intervention.

Grouping objects
Use the Object group editor to organize situations, managed systems, and historical configurations into named collections that can be applied with a single action.

Navigator views for logical organization
Create custom Navigator views to organize your monitored environment into logical hierarchical groupings. You can also create Navigator views and build workspaces that are targeted for specific points of view, such as executive, operator, or subject matter expert for an application domain.

Tivoli Enterprise Portal tour

Welcome to the Tivoli Enterprise Portal tour. In under 10 minutes, this hands-on practice introduces you to some of the major features.

Navigator

The Navigator Physical view shows the physical hierarchy of your network, with *enterprise* at the top, followed by the *operating platform*, and so on. Have a look at the levels of the Navigator:

1. Open the operating platform level (Linux Systems, UNIX Systems, Windows Systems, or z/OS Systems) by clicking ⊞ Expand to the left of its icon.

   ⬛ Enterprise
       ⊞ 📁 Windows Systems

   Opening a level in the Navigator reveals the next level in that branch.

2. Open the operating platform level to see the names for the *systems* running on that platform.

   ⬛ Enterprise
       ⊟ 📁 Windows Systems
           ⊞ 🖥 JDOE

3. Open any system to see the monitoring *agents* installed on that system for monitoring applications and resources, and, below agents, the elements, or *attributes*, for which the agent can collect data for reporting and testing.

    ▦ Enterprise
        ⊟ 📁 Windows Systems
            ⊟ 🖥 JDOE
                ⊟ ◈ Windows
                    🖵 Disk
                    🖵 Monitoring Services
                    🖵 Memory
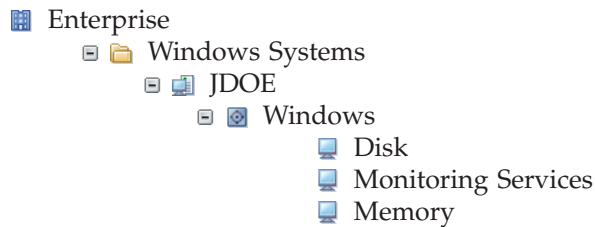
⊞ Expand and ⊟ Collapse enable you to open and close levels in the Navigator without selecting an item.

Tip: You can close the tree entirely by double-clicking the ▦ Enterprise item. To expand the tree, double-click again.

Workspaces

Every item in the Navigator has a default workspace that opens when you select it. Multiple workspaces can also be accessed from a single navigator item. After you log into the Tivoli Enterprise Portal Server, the Tivoli Enterprise Portal window opens. The top item in the Navigator, ▦ Enterprise, is selected and its workspace is displayed.

- Select another Navigator item by clicking the 📁 icon for the operating platform or the name.

The workspace for the operating platform you selected replaces the one that was previously displayed. The workspace displayed when you click a Navigator item is the *default workspace* for that item. Your IBM Tivoli Monitoring product might come with multiple workspaces for some items, and you can create additional workspaces of your own.

- Click ▦ **Workspace gallery** to see thumbnail graphics of all the workspaces that you can open for this Navigator item.

If this is a new installation, you will see the splash screen instead of a thumbnail version of the workspace in the gallery until you or another user who is logged on to the same Tivoli Enterprise Portal Server opens the workspace for the first time. A ✔ check mark by a workspace name indicates that it is the default workspace for this Navigator item. You can reorder the workspaces in your copy of the gallery by clicking a thumbnail graphic and dragging it on top of the workspace to swap positions with.

Together, the Navigator and workspaces show the facets of your Tivoli managed enterprise from the highest level to the most detailed.

Views

Tivoli Enterprise Portal offers many different views that you can add to your workspaces.

The view title bar has buttons for creating another view by splitting the view ⊟ horizontally or ⫼ vertically, ☐ for maximizing the view, and for ✖ deleting the view.

To add a different view complete the following steps:

1. Select a view, then click ▱ **Split Horizontally**.

   This tool divides the view space in half horizontally to create a copy of the original view.

   If you cannot split a view and the tools for the notepad and other views are disabled (dimmed), your user ID does not have Workspace Author Mode permission.

2. Click your desired view type from the toolbar. For example click 📝 **Notepad**

   When you click a view tool, the mouse pointer changes to an icon of the chosen view type (hand icon on Linux).

3. Click inside the view pane. The new view replaces the previous view.

To rearrange views complete the following steps:

1. Drag the title bar of one view and drop over another view.

   While dragging the title bar, you should see a semi-transparent copy of the view.

2. Release the mouse button to switch the views.

Query-based views

> The table and chart views are the first step to getting meaningful information from the data being collected. When you understand what values and states are causing problems, you can refine your views to show what is important. For attribute groups with attributes that can be expressed meaningfully as objects and show their relationship to other attributes in the same group, there is also the *Relational Table Based* topology view.

> 🖽 The table view and 🌐 📊 🖾 🖾 🕐 📊 chart views display data the monitoring agents have gathered from the systems where they are running.

> 🔀 The topology shows the arrangement of monitored components associated with its Navigator item.

Event views

> The message log view and console views (universal message, situation event, common event, and Tivoli Enterprise Console) update automatically to show new events as they arrive and changes in event status as they occur. The graphic view gives you a pictorial alternative to the Navigator for indicating alerts.

> 📋 The message log view shows the status of events that have been opened on the entire monitored enterprise and can include up to 100 row entries at a time.

> 📃 The universal message console view displays situation and policy activities as they happen, such as when a situation has been created or deleted or a policy has been activated.

> 🖼 The graphic view places Navigator items and their alerts as icons on a map or a picture to represent your monitored environment.

> 🖽 The situation event console view shows the status of situation events that have been opened on this branch of the Navigator. The console has a toolbar for quick filtering and a menu for responding to alerts.

> 🔳 The common event console view integrates events from multiple event repositories: Tivoli Enterprise Monitoring Server, Tivoli Enterprise Console

event server, and Netcool/OMNIbus ObjectServer. This view has the same functional capabilities as the situation event console view, such as the ability to sort and filter events.

The Tivoli Enterprise Console event viewer can integrate events from the Tivoli Enterprise Console Server with those from the Tivoli Enterprise Monitoring Server.

Other views

Tivoli Enterprise Portal workspaces can also have any of the following specialized views:

The notepad view opens a simple text editor for writing notes or other text that can be saved with the workspace.

The take action view enables you to send a command to a managed system.

The terminal view starts a 3270 or 5250 session for working with z/OS applications, or a Telnet session for working with the TCP/IP network.

The browser view opens the integrated browser for accessing web pages.

The graphic view places Navigator items and their alerts as icons on a map or a picture to represent your monitored environment. Alerts show on these icons just as they do in the Navigator.

Situations

In addition to providing a map of your enterprise, the Navigator can alert you of changing conditions in the systems you are monitoring. When a condition changes, the associated item is marked with an icon: ● Fatal, ● Critical, ▲ Minor, ▲ Warning, ▣ Harmless, ▣ Informational, or ● Unknown. The Navigator places one of these icons, called *alert indicator* or *event indicator*, at every level of the hierarchy so you can see an alert even if a Navigator branch is closed.

If you see any of these icons in the Navigator now, move the mouse pointer over an event indicator to open a flyover listing of the situations that caused the events at that level of the Navigator and below.

Tivoli Enterprise Portal runs tests called *situations* on systems where monitoring agents are installed. When the conditions of a situation have been met, an event occurs and an event indicator is displayed over the applicable items in the Navigator.

Open a situation using the following steps:

1. Right-click a system name in the Navigator and click 📁 **Situations**. The Situation editor opens with a list of situations on the left. This is where you can view and edit situations or create new ones.
2. Click a situation name in the tree to see the Situation editor tabs. *fx* Formula to view, add, and change the conditions being tested; 📁 Distribution to view and specify the systems on which to run the situation; ● Expert Advice to write comments or instructions to be read in the event results workspace; ⚡ Action to send a command to the managed system; and ◔ **Until** to close the event after a period of time or when another situation is true.
3. Click **Cancel** to close the Situation editor.

You cannot open the Situation editor if your user ID does not have permission to View Situations.

Properties

In this last part of the tour, open the Properties editor, where you can edit the appearance and content of the views chosen for a workspace.

1. Click ▦ **Properties**. The Properties editor opens with a list of the views on the left, grouped by type. The general workspace properties appear on the right.

2. Select one of the views in the **Properties** tree. All types of views have a 𝐴𝑞 **Style** tab, with different options for the type of view. The table, charts, and topology views have 🅹 **Query** and ⚒ **Filters** tabs, table and topology have a ▤ **Thresholds** tab, and terminal emulators have a 📄 **Configuration** tab. Whenever you want to change the data elements that display for a chart or table or change the styling of a view, go to the Properties editor.

3. Click **Cancel** to close the Properties editor.

Conclusion

You now know the basics on the main functionality of the Tivoli Enterprise Portal. For practice, try splitting this view and apply a ▦ Browser view.

Chapter 2. Using the Navigator

The Navigator provides a hierarchical view of your enterprise. At the highest level, you can get a high-level overview of the status of your monitored environment. From there, you can navigate to specific monitored resources to check activity and investigate problems.

At every level, event indicators alert you to changes in system or application conditions.

Navigator overview

The Navigator provides a hierarchical view of your enterprise. At the highest level you can get a high level overview of the status of your monitored environment. From there you can navigate to specific monitored resources to check activity and investigate problems.

Types of Navigator views

Navigator Physical view

The default Navigator view is Physical and shows the hierarchy of your monitored enterprise. The Navigator Physical view is a discovered view: as new managed systems come online or when they become disconnected, the view is adjusted accordingly.

Enterprise, or root, level contains all managed systems
   Platform is the operating system
      System, or node, is the name of the computer or z/OS image
         Agent type. Multiple instance agents have child items; some are grouped into one folder
            Resource category being monitored

Here is an example of a system named JDOE, running on Windows, with the Windows OS agent.

Enterprise
   Windows Systems
      JDOE
         Windows OS
            Disk
            Monitoring Services
            Memory

You can collapse the tree by double-clicking the Enterprise item, and expand it by double-clicking again.

Custom Navigator views

Your Tivoli Enterprise Portal configuration can also have custom-designed Navigator views. These views are selectable from the Navigator toolbar, navigable and show event indicators (described below) in the same way as the Navigator Physical view. Unlike the Navigator Physical view, custom Navigator views can be edited. You can, for example, design a Navigator view for Manufacturing and another for Marketing.

Your user ID can be assigned to all or a subset of the defined Navigator views, and your access in any of these views can be restricted to a certain branch. To see the list of available Navigator views, click the ▾ list box in the Navigator toolbar.

A ✛ small white cross over a Navigator item icon means one or more other Navigator items share the characteristics of that item. This happens when an item from one Navigator view has been copied to another Navigator view. All the managed systems, workspaces, link definitions, situations, and policies associated with the original (source) item are applied to the new item. And from then on, changes to one item (such as a new workspace is added) are applied to the other. Right-click a Navigator item and click ⚏ **Show Navigator List** to see the other Navigator views it is contained in.

Navigator Logical view

Users initially have one custom Navigator view called *Logical* with a single Navigator item for 📁 Enterprise. Click the ▾ list box and select the Navigator Logical view from the list.

Workspaces

The *workspace* is the working area of the application window. Selecting an item in the Navigator opens its default workspace. The workspace contains views of information for that item in the Navigator.

If the item has multiple workspaces created for it, you can click ▦ **Workspace Gallery** to see thumbnail images of them and select another workspace to open. You can also select any additional workspaces associated with a navigator item by right-clicking on the navigator item and selecting the **Workspaces** menu option, or by selecting that same option from the **View** menu found on the main menu bar.

You can move around the Navigator without selecting workspaces by using the vertical scroll bars and clicking ⊞ to expand a list or ⊟ to collapse it.

Use ⬅ and ➡ to return to previously viewed workspaces in the order they were initially visited. The data in those workspaces is not cached, so any chart or table views are refreshed with the most recent data samplings when you revisit them.

Navigator tools

Use the tools in the Navigator view toolbar to update the display when managed systems have come online or gone offline, to open another Navigator view, to edit custom Navigator views, and to collapse the view.

⟳ 2 Apply pending updates

The **Apply pending updates tool** updates the Navigator tree with the correct status of agents, in response to the count of agents that have been removed from or added to the Tivoli Enterprise Monitoring Server. The counter is displayed next to the tool and in the status bar. Refresh is enabled when the Navigator view has been edited, such as a new item has been added, so that you can update the view with the changes. This tool is dimmed when no updates are required. (See "Refreshing the Navigator" on page 29.)

▨ 6 Application support event

After you start the Tivoli Enterprise Portal, the application support version is read from the portal client and the portal server. Any discrepancies in versions between client and server are noted in the Application Support Mismatch list. (See "Responding to an application support event" on page 30.)

🔍 Find Navigator item

Use 🔍 **Find** in the Navigator toolbar to look for a Navigator item by such criteria as managed system name or associated situation name. When the search has completed, any items that match the criteria are listed in the **Find results** table where you can switch to the default workspace for a found item or open it in a new window.

✎ Edit navigator view

Opens the Navigator editor for creating, deleting, and modifying Navigator views. (See Chapter 9, "Navigator views for logical organization," on page 237.)

If the tool is disabled (dimmed), your user ID does not have Modify permission for Custom Navigator Views.

▾　Opens a list that you can select a Navigator view from. The views available are those assigned to your user ID. When you select another Navigator view, a ⛛ tab is added to the bottom of the Navigator so you can quickly switch among the views by clicking the tab. (See "Opening a Navigator view" on page 28.)

◀　Minimizes the Navigator and expands the adjacent workspace views to fill the gap. Click ▶ to restore its original size. You can also click ⛛ *<Navigator name>* in the right-most section of the status bar to restore the Navigator view; or right-click it to see a list of Navigator views to select from.

⊟　Splits the Navigator horizontally to create a new workspace view. (See Chapter 6, "Customizing workspace views," on page 85.)

If the tool is disabled your user ID does not have Author Mode permission.

⊞　Splits the Navigator vertically to create a new workspace view. (See Chapter 6, "Customizing workspace views," on page 85.)

If the tool is disabled your user ID does not have Author Mode permission.

Situation event indicators

When you see a small colored icon overlay the Navigator icons, you are seeing an *event indicator* . An event indicator is displayed when a situation (a test of certain conditions) running on a system becomes true.

These indicators alert you that conditions have changed and need attention.

Event indicators can be, from highest severity to lowest, ⬤ Fatal, ⊗ Critical, ⚠ Minor, ⚠ Warning, ▣ Harmless, ▣ Informational, or ⬤ Unknown. As you move up the Navigator hierarchy, multiple events are consolidated to show only the indicator for the event with the highest severity. Go to the lowest level of the hierarchy to see individual event indicators over attribute groups with values that have met situation definitions.

Move the mouse pointer over an indicator to open a flyover listing of the situations that are true for the Navigator item and any branching items. This flyover list might also show a status icon on the right:

☑ Acknowledged	The situation event has been acknowledged.
ⓘ Expired	The acknowledgement has expired and the situation is still true.
↺ Reopened	The acknowledgement was canceled before it had expired and the situation is still true.
⬤ Stopped	The situation has been stopped.
❗ Error	The situation is not functioning properly.
⬤ Status Unknown	The monitoring server detects that an agent is offline. The agent might have been taken offline intentionally, there might be a communication problem, or the agent or the system it is running on might have stopped or be failing. The situation flyover listing on this icon shows *STATUS_UNKNOWN, which is not actually a situation, but the notation for a problem on the managed system. Consider recycling the agent, described in "Managing monitoring agents" on page 400.

More... indicators

You might sometimes see ⋱ **More** indicators in the Navigator. These indicators are used to keep the Navigator tree compact so that you can see more in the viewable area without having to scroll.

Click ⋱ **More** to open that branch of the tree.

When clicking ⊞ Expand on a Navigator branch of more than, by default, 25 child items, you are prompted to enter the number to expand at one time. You can also right-click a Navigator item and click ⊞ **Expand** to open the **Expand child items** window. After you enter the number to expand, a group of child items opens. Click ⋱ **More** to open the next group of child items.

In large managed environments this feature helps you control the Navigator display and its performance. In this example, the expansion limit of 4 was chosen for the Windows Server branch:

    ⊟ 📁 MYCOMPUTER

        ⊟ ◉ Windows Server

- Disk
- Enterprise Services
- Memory
- Network
- More... 4 : 8

Expanding the Navigator in increments

More indicators in the Navigator keep the tree compact. Click the indicator to open that branch of the tree. You can specify how many items to expand at one time.

About this task

When clicking ⊞ Expand on a Navigator branch of more than, by default, 25 child items, you are prompted to enter the number to expand at one time.

Procedure

1. Set the number of items to expand:
   a. If the Expand child items window is not open, right-click the Navigator item whose child items you want to expand incrementally and click ⊞ **Expand**.
   b. Enter the number to expand and click **OK** to open that number of child items.
   c. Click ⚡ **More** when you want to open the next group of child items.
2. Expand child items. You can expand all the child items of a Navigator item. The entire branch does not expand, just the children of the selected item. If the Expand child items window is not open, right-click the Navigator item whose child items you want to expand and click ⊞ **Expand Child Items**.

Example

In large managed environments the expand feature helps you control the Navigator display and its performance. In this example, the expansion limit of 4 was chosen for the Windows Server child items:

- Enterprise
  - Windows Systems
    - Computer1
      - Windows Server
        - Disk
        - Enterprise Services
        - Memory
        - Network
        - More... 4 : 8

This example illustrates the result of selecting ⊞ **Expand Child Items** at the platform level of the Navigator. The site is monitoring 500 Windows-based systems. The user can quickly see which systems have situation events and at what severity.

- Enterprise
  - Windows Systems

      ⊞ 📁 Computer1
      ⊞ 📁 Computer2
      ⊞ 📁 Computer3
      ⊞ 📁 Computer4
      ⊞ 📁 Computer5
      ⊞ 📁 Computer6
      ⊞ 📁 Computer7
      ⊞ 📁 Computer8
      ⊞ 📁 Computer9
      ⊞ 📁 Computer10
      ⊞ 📁 Computer11
      ⊞ 📁 Computer12
      ⊞ 📁 Computer13
      ⊞ 📁 Computer14
      ⊞ 📁 Computer15
.
.
.
      ⊞ 📁 Computer500

Logical and custom Navigator views

    A custom Navigator view can be designed to show the same managed
    system more than once. If you have a managed system repeated
    somewhere else in the Navigator view, when you use the expand feature
    on one Navigator item for that managed system, any other items for that
    managed system will expand at the same time.

Collapsing and expanding the Navigator

To provide more space for workspace views, you can hide the Navigator, then
restore it when you want to see it.

You can also can save the workspace with the Navigator minimized.

Procedure

* To collapse the Navigator view, click the ◀ **Collapse** bar on the right border.
  (You can also drag the right border if you prefer to manually adjust the
  Navigator width.) The adjacent workspace views expand to fill the gap and the
  status bar shows the Navigator view icon and name in the right-most section.

* To restore the Navigator view, either click the ▶ **Expand** bar on the left border or
  click ⌗ *<Navigator name>* in the status bar. There is an additional option to open
  a different Navigator view instead of restoring the original by right-clicking ⌗
  <Navigator name> to see and select from a list of available Navigator views.

Finding Navigator items

The Navigator shows all the managed systems that run on an operating system
and this list can get very long. Use the Find tool to locate Navigator items based
on properties such as IP address or associated situations.

About this task

Take these steps to find a Navigator item:

Procedure

1. Click 🔍 **Find** in the Navigator toolbar.
2. If the item you want to search for is in a different Navigator view, select it from the **For Navigator** ▾ list.
3. Enter the values of the properties to search by:

Properties	Value
IP Address	(Distributed systems only) If you know the IP address, you can enter it here. Example: SCAN == 9.55.1* finds computers that start with 9.55.1, such as 9.55.105.180 and 9.55.110.180
Hostname	(Distributed systems only) Specify all or part of the host name here. Example: SCAN == MYHOST?? finds computers that start with MYHOST and have two trailing characters, such as MYHOSTEL and MYHOST33 but not MYHOSTA or myhostess.
SMFID Name	(z/OS-based systems only) Find the item by the session ID value.
Sysplex Name	(z/OS-based systems only) Specify the item by system complex name.
Product Code	Select the product name from the ▾ list to specify the managed system type to find.
Managed System Name	The managed system name is often the same as the host name. Example: SCAN == *UAGENT0? finds Universal Agents Navigator items such as onionASFSdp:UAGENT00 and garlicASFSdp:UAGENT05.
Associated Situation Name	Enter a situation name to find the Navigator items that the situation is associated with. (For an alert indicator to display in the Navigator for a situation that has become true, the situation must be associated with a Navigator item.)

4. If you want to show the found Navigator items at the agent level (and subagent if applicable) but to exclude the child Navigator items (attribute level), select the ☐ **Include only managed system items** check box.
5. If you want to control the number of Navigator items to list per page in the find results, click **Advanced** and specify the page size. The default is 100 rows per page. By reducing the number, you can have the results displayed more quickly and in smaller, more manageable pages; then move through the pages with ⬇ Page down and ⬆ Page up. By increasing the number, you increase the processing time but you can see all the results on one table page.
6. Click **Find**. The **Find results** area is appended to the window with a list of Navigator items that meet the criteria.
7. Click a Navigator item in the Find results area and click **Switch to** after selecting any or none of these check boxes:

    ☐ **Open workspace in new window**
    > Keep the current workspace open and open the workspace at the selected Navigator item in a new window (or separate tab if you are in browser mode and your browser is configured for tabbed windows).

□ **Close after switching to workspace**

Close the Find window after opening the workspace for the selected Navigator item.

Results

The default workspace for the selected Navigator item is opened.

Example

In this example, the find criteria looks on the MySystem computers for Navigator items with associated situations that have "CPU" in the situation name.

Find criteria

Properties	Value
IP Address	
Hostname	abc == 'MySystem?'
SMFID Name	
Sysplex Name	
Product Code	
Managed System Name	
Associated Situation Name	abc == '*CPU*'

Find results

Item	View	Item path
Process	Physical	WindowsOS/MySystem1/Windows Systems/Enterprise
Processor	Physical	WindowsOS/MySystem1/Windows Systems/Enterprise
Process	Physical	WindowsOS/MySystem2/Windows Systems/Enterprise
Processor	Physical	WindowsOS/MySystem2/Windows Systems/Enterprise

What to do next

You can keep the window open to switch between one found Navigator item and another or to change the criteria and find other Navigator items.

Opening a Navigator view

The initial Tivoli Enterprise Portal workplace has a Navigator Physical view and a Logical view and any other views that were predefined for your product. You can switch between Navigator views to see the items and workspaces of another Navigator.

About this task

Depending on whether the Navigator is displayed or hidden, take either of these steps to open a different Navigator view:

Procedure

- Click the View ⏷ list in the Navigator toolbar and select the Navigator view you want to open.
- If the Navigator is not displayed because it is collapsed, you can right-click ⛕ *<Navigator name>* in the right-most section of the status bar to see and select from a list of available Navigator views.

What to do next

If you do not see the Navigator view you are looking for, it is possible your user ID does not have the authority to access this Navigator view. Ask your IBM Tivoli Monitoring administrator.

If there are many Navigator views in the list, you can have them reordered through the **Navigator Views** tab of the Administer Users window.

Related tasks:

"Adding a user ID" on page 369
Create a user ID for all users that should be able to log onto the Tivoli Enterprise Portal Server using a portal client or the `tacmd tepsLogin` command. A user ID is also required for IBM Dashboard Application Services Hub users who request monitoring data. You can use the default user profile or copy the profile of an existing user.

Refreshing the Navigator

The Navigator refreshes automatically whenever the status of agents changes between offline and online. The Navigator does not update automatically when an agent is removed from or a new agent is added to the managed network. When the tool is enabled, click ⟳ **Apply pending updates** to manually refresh the Navigator.

Before you begin

If an agent name is dimmed (grayed out), the agent has stopped or is not responding. It might also indicate that Tivoli Enterprise Portal has timed out before detecting the availability of the agent.

For custom Navigator views, ⟳ **Apply pending updates** is enabled in the Navigator toolbar whenever you have edited the view in the Navigator editor (added, removed, or renamed items). The tool is not enabled in the Navigator view when an agent has been removed from the Tivoli Enterprise Monitoring Server. In these cases, the agent appears to be offline. To remove this agent from the Navigator display, delete the Navigator item through the Navigator editor.

About this task

A count of changes to the Navigator view is kept as agents are removed from, disabled, or added to the monitoring server. The counter shows as a superscripted number next to the tool (such as $\curvearrowright^9$ for nine updates) and in the Navigator view status bar. You must refresh the Navigator to show these changes. Take one of the following steps:

Procedure

- Click ⚙ **Apply pending updates** in the Navigator toolbar.
- Click ⚙ **Navigator update pending** in the Navigator status bar.

Results

The agent status is updated in the Navigator. (You might not see any change if the branch containing the agent is ⊞ collapsed.) The tool is disabled until the another status change is detected, at which time the counter starts again. If the workspace has automatic refresh turned on, the Navigator is updated at the end of the refresh interval.

Responding to an application support event

Use the **Show new or updated application support** tool to see a count of the version mismatch messages, to view a list of application support version mismatches between the Tivoli Enterprise Portal client and server, and to see instructions for correcting the discrepancy.

The number that is displayed after the tool, such as 📋 [8] for eight mismatch events, is a count of the events that have accumulated.

Procedure

1. When the 📋 **Application support event** tool is enabled in the Navigator toolbar, click it. The dialog box shows the list of mismatched applications and their support versions. Each event shows a Warning or Informational indicator followed by the name of the monitoring agent, the version of application support applied at the portal client, and the version applied at the portal server. No portal client version information is available when a new monitoring application is initially installed.
2. Complete the instructions that are displayed after the mismatch event list:

   Browser client and Java WebStart client
   > If you are running the browser client or Java WebStart, any required application support updates are downloaded automatically after you log on to the portal server.
   >
   > ⚠ *An application support version mismatch has been detected. Click Help for more information on possible causes and actions.* Warning events indicate that the mismatch was found at client startup.
   >
   > ℹ *New or updated application support is available. Restart the Tivoli Enterprise Portal client to update your application support.* If application support updates become available while you are logged on, an Informational event occurs to show that application support is waiting to be applied. Log out and restart the client to have application support updates applied.

   Desktop client
   > Application support updates to the desktop client must be made by running the installer on the Tivoli Monitoring Agent installation media.
   >
   > ⚠ *An application support version mismatch has been detected. Use the Tivoli Monitoring installer to add the associated application support to the portal desktop client.* Warning events indicate that the mismatch was found at client startup.

New or updated application support is available. Use the Tivoli Monitoring installer to add the associated application support to the portal desktop client. Informational events indicate that the mismatch was found after logon to the portal server.

Results

After logging off and logging onto the Tivoli Enterprise Portal, the **Application support event** tool is disabled. If it is not, contact your network administrator or IBM Support.

The first time the **Show new or updated application support** tool is enabled and after a new event arrives, a "New" indicator is displayed on the tool: . After you open and then close the dialog box, the indicator is no longer shown: . The indicator provides a convenient visual cue when new events arrive that you have not yet seen.

Usage note:

If the client application support version is greater than the portal server's application support version, there is no need to recycle the client. The new support will work with the earlier support that is present on the portal server. Recycling the client does not collect the earlier version of support from the portal server unless the Java cache has been cleared to remove the newer version of application support. To clear the Java cache complete the following steps:

1. If Tivoli Enterprise Portal is running, exit by closing the browser window.
2. Start the Java Control Panel. Go to **Start > Settings > Control Panel**, then double-click the Java Control Panel icon.
3. In the Java Control Panel window, on the **General** tab delete the temporary internet files. When the JAR cache is cleared, click **OK**.
4. Restart the Tivoli Enterprise Portal client.

Removing an event item

When you open an event results workspace, a new item is displayed in the Navigator. If you use event acknowledgement, other icons overlay the situation icon. If you close the event, all overlay icons are removed, but the Navigator item itself is not removed until you end your Tivoli Enterprise Portal session. If you do not want to have the event item in the Navigator, you can remove it.

This does not affect the event or operation of the situation, only its appearance in the Navigator.

Procedure

To remove an event item, in the Navigator right-click the situation name and click **Remove**. The event item is removed from the Navigator. To see it again, you need open its workspace by selecting the true situation from the event flyover list.

Example

Here are some examples of event icons that you might see in the Navigator:

A critical event.

	A warning event.
	A warning event that has been acknowledged.
	A warning event whose acknowledgement has expired.
	A warning event whose acknowledgement was removed while the event was still open.
	A true situation whose event has been closed.
	An agent whose status is unknown. The monitoring server detects that the agent is offline. It might have been taken offline intentionally, there might be a communication problem, or the agent or the system it is running on might have stopped or be failing.

Removing an offline managed system Navigator item

Use the Managed System Status workspace to clear offline managed systems from the Navigator view.

About this task

When a monitoring agent goes offline or its host name gets changed, its branch is dimmed in the Navigator, as in the following example, and its status is shown as *OFFLINE in any table view that uses the Current Managed Systems status query.

- Enterprise
    - Windows Systems
        - MyComputer
            - Universal Agent
            - Windows OS
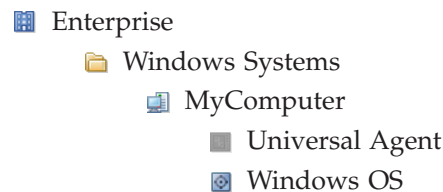
The Enterprise Navigator item has a Managed System Status workspace that you can use to check for any offline managed systems and remove them from the Navigator view. When the monitoring agent comes online again, it will reappear. Use these steps to clear an offline entry:

Procedure

1. Click the Enterprise Navigator item. The default workspace opens.
2. Click **Workspace gallery**, scroll to **Managed System Status**, and click the workspace.
3. Right-click an ***OFFLINE** row, then click **Clear offline entry**.

Results

The offline entry is now cleared from the Navigator and will not appear again until the managed system comes online.

What to do next

If this is a custom Navigator view with only one managed system, it is possible for the item to continue to appear in the view. This is possible when the item was added using drag-and-drop in the Navigator editor. Use the Navigator editor to manually remove the item, described in Deleting a Navigator view or item. If the managed system comes online again, it does not appear in the custom Navigator view. If you want it to appear there, assign the managed system in the Navigator item properties as described in Navigator item properties.

Related tasks:

"Removing a managed system and monitoring agent" on page 409
Use ⊠ **Remove** to stop a managed system and permanently remove it from the
Navigator Physical view.

Chapter 3. Using workspaces

The workspace is the working area of the Tivoli Enterprise Portal window, divided into panes to show different types of views. You can start monitoring activity and system status immediately with the predefined workspaces.

You can tailor your own workspaces to give you summary overviews or to look at specific conditions.

Workspace characteristics

Every Navigator item has at least one predefined workspace that you can open. Every workspace characteristics such as editable properties and views.

Views

A view is a windowpane, or frame, in the workspace containing a chart or table showing data from one or more monitoring agents. Other types of views such as the topology view and graphic view can give a broader overview of the network. Specialized view such as the browser view and terminal view are also available. You can increase the number of views in a workspace by splitting a view into two separate views.

The data for a table, chart, or relational table-based topology view is chosen by the query it uses. Collectively, they are called *query-based views*. The query specifies the attributes to include in the view. Although each view uses one query, you can add more views to the workspace, and each can use a different query. The queries can be for different monitoring agents, including those for the Tivoli Enterprise Monitoring Server for showing information that is common to your monitored environment (such as all the managed systems and all the situation events). You can also include queries of JDBC or ODBC data sources by writing custom SQL queries.

Properties

Every workspace has a set of properties associated with it: general properties that apply to the entire workspace, and properties for each view in the workspace. Use the Properties editor to customize the workspace characteristics and to change the style and content of each view.

You can also keep the original workspace intact and create another workspace for the same item in the Navigator, customizing it for the types of views you want and the information reported in charts and tables.

Changes you make to a workspace are available only to your user ID. System administrators can work in Administration mode to create and edit workspaces that will be available to all users on the managed network.

Links

The link feature enables you to define a link from one workspace to another. Then you can quickly jump to a related or more detailed workspace to investigate system conditions.

The simplest type of a link originates from the Navigator item: When you right-click that Navigator item, the pop-up menu shows the defined links for the item. Select one to open the linked workspace.

A more specific link originates from a table or from a chart data point to another workspace. Information from one of the attributes in the selected row, bar, pie segment, or plot point is used to determine the content of the target workspace.

You can also define more complex links and use the predefined links that come with your IBM Tivoli Monitoring product.

Navigator level

The monitoring agents available for reporting in a workspace are those assigned to that branch of the Navigator. If you are not sure which monitoring agents are included, do one of the following:
- Expand the branch of the Navigator
- Right-click the Navigator item and select Properties to see which managed systems are assigned.
- Open one of the workspaces at the enterprise, platform, or system level of the Navigator Physical view

This same principle applies to attribute groups. The lowest level of the Navigator Physical view, for example, is the attribute level. The views you can show for the workspaces at that level can draw only from the attribute groups represented by that level. If you were to build a workspace for the 🖥 **Disk** Navigator item, for example, you could create a chart with data from the **Logical Disk** attributes and another with data from the **Physical Disk** attributes.

Organization of predefined workspaces

The Enterprise Navigator item has workspaces that query the Tivoli Enterprise Monitoring Server. Use these predefined workspaces to get status information about the monitoring server and monitoring agents and about situations and policies.

Use the ▦ **Workspace Gallery** to see what is available for the Navigator item. These are the Enterprise Navigator item workspaces and the workspaces they link to:

Enterprise Status
> The default workspace is Enterprise Status, which gives an overview of the situation event status throughout your enterprise.
> - Event Details - Similar by Situation Name
> - Event Details - Similar by Source
> - Event Details - Similar by Resource

Manage Tivoli Enterprise Monitoring Servers
> The Manage Tivoli Enterprise Monitoring Servers workspaces provide a visual health check of the monitoring servers in your enterprise and the application support that has been applied.
> - Installed Catalogs - Enterprise View
> - Installed Catalogs - Remote Server
> - Protocols
> - Situation Status
> - System Information

Managed System Status

The Managed System Status is a list of monitoring agents in your managed network and their ONLINE or OFFLINE status. The linked workspaces are only available for online managed systems.

- Audit Log
- Agent Operations Log
- History Exports

EIB Change Log

This workspaces displays entries in the Enterprise Information Base (EIB) log. The EIB is a database used by the Tivoli Enterprise Monitoring Server to store situation, policy, user definitions, and configuration information.

Self-Monitoring Topology

The Self-Monitoring Topology workspace provides a high level overview of your managed infrastructure and its health.

Deploy Depot Package List

The Deploy Depot Package List workspace shows the installation packages that are available in the agent depot.

Deployment Status Summary

The Deployment Status Summary workspace shows summary status information about remote agent deployments.

- Deployment Status by Deploy Group
- Deployment Status by Product

Deployment Status Summary by Transaction

The Deployment Status Summary by Transaction workspace shows summary status information about remote agent deployments, sorted by transaction.

Related reference:

"Enterprise Status workspace" on page 57
At the highest level of the Navigator Physical view is the ▦ Enterprise item. The default workspace for the Enterprise is called Enterprise Status.

"Manage Tivoli Enterprise Monitoring Servers workspaces" on page 398
Use the Manage Tivoli Enterprise Monitoring Servers workspaces and situations to gather information for diagnosing and correcting typical monitoring server configuration issues.

"Managed System Status workspace" on page 392
The Managed System Status workspace is one of the predefined workspaces for the Enterprise Navigator item, and offers a high level overview of your managed network. The workspace has a single view: *Managed System Status table view*.

"EIB Change Log workspace" on page 396
The Enterprise Information Base (EIB) is a database used by the Tivoli Enterprise Monitoring Server to store situation, policy, user definitions and configuration information. The EIB Change Log view is a log with details of changes made to information in the database.

"Self-Monitoring Topology workspace" on page 397
The Self-Monitoring Topology workspace provides a high level overview of your managed infrastructure and its health and introduces you to the Tivoli Enterprise Portal self-monitoring capabilities.

"Agent deployment workspaces" on page 400
Use the agent deployment workspaces to see your package list, the results of a deployment prerequisite check, and the progress of the agent deployment process.

Opening a workspace

Use the Navigator to open the default workspace for the selected item, then the Workspace Gallery to see and select from a thumbnail display of available workspaces.

Before you begin

As well as a convenient way to open workspaces, you can also move the workspaces around the gallery to change their order. The changes you make are saved with the workspace definition for the Navigator item. If you are in workspace administration mode when you reorder the thumbnail graphics, the reorganization will be reflected in the workspace gallery of all Tivoli Enterprise Portal clients connected to this Tivoli Enterprise Portal Server.

About this task

Every item in the Navigator has a workspace associated with it, called the *default workspace*. Some items have multiple workspaces that you can open, although only one workspace can be open in the Tivoli Enterprise Portal window at one time. Some workspaces are only accessible by linking to them from another workspace.

Procedure

- To open the default workspace, click the item name or icon in the Navigator:
    - Physical view: ▦ , 🗀 , 🖳 , 🔘 , or 🖥 .
    - Logical view or other custom Navigator view: 🗀 .

  The default workspace for that item is displayed, replacing the workspace of the previously selected item. If the workspace shows no data for a chart or table

view, it means there is no data to display. This can occur with monitoring data that is not constantly generated, such as Archive Errors, which collects data only when archive errors occur.

- To open another workspace associated with the Navigator item:
  1. Click ⊞ **Workspace gallery**.
  2. Click the thumbnail graphic of the workspace to open. (To open the workspace in a new window instead of replacing the current workspace, use Ctrl + Shift + click the thumbnail graphic.)

If this is a new installation, you will see the splash screen instead of a thumbnail version of the workspace in the gallery until you or another user who is logged on to the same Tivoli Enterprise Portal Server opens the workspace for the first time. A ✓ check mark by a workspace name indicates that it is the default workspace for this Navigator item. You can reorder the workspaces in your copy of the gallery by clicking a thumbnail graphic and dragging it on top of the workspace to swap positions with.

Related tasks:

Starting workspace administration mode
Typically, workspaces you create are saved with your user ID and no other users can see them. You can turn on administration mode so that all workspaces you edit and save while in this mode are shared with all users who log on to the same Tivoli Enterprise Portal Server.

"Opening a new window"
Have multiple workspaces open on your desktop at the same time by opening multiple Tivoli Enterprise Portal windows.

Related reference:

"Tabbed workspaces" on page 40
Use the tabbed pages capability of your browser to open workspaces, linked workspaces, and Navigator views in new tabs.

Opening a new window

Have multiple workspaces open on your desktop at the same time by opening multiple Tivoli Enterprise Portal windows.

About this task

You can add more dashboard views to your desktop by opening more Tivoli Enterprise Portal windows. Any changes you make these window are saved with the application.

Procedure

- To open a new window in the desktop client, click 🗋 **New Window**.
- To open a new window in the browser client running in Internet Explorer, press Ctrl + N.
- To open another workspace in a new window and keep the original intact in this window, click ⊞ **Workspace gallery** and Ctrl + Shift + click the workspace.

Results

The new window is opened as a duplicate of the original; any changes you make to the new window are independent of the original.

What to do next

Any previously visited workspaces are retained from the parent window; use ⬅ and ➡ to revisit them. Further navigation to other workspaces in either window, however, is independent of the other window.

You can close duplicate windows (click **File > Close)** or the original; the work session remains active as long as one window is open.

Related tasks:

"Opening a workspace" on page 38
Use the Navigator to open the default workspace for the selected item, then the Workspace Gallery to see and select from a thumbnail display of available workspaces.

Related reference:

"Tabbed workspaces"
Use the tabbed pages capability of your browser to open workspaces, linked workspaces, and Navigator views in new tabs.

Tabbed workspaces

Use the tabbed pages capability of your browser to open workspaces, linked workspaces, and Navigator views in new tabs.

Browser client and browser settings

When your browser supports tabbed web pages, the Tivoli Enterprise Portal browser client uses the browser's tab settings to determine how to open a workspace: When tabs are enabled, the workspace is opened in a new tab. You can set the properties of a workspace or the target of a workspace link to always open in a new tab, or you can open a workspace in a new tab by holding down the Ctrl + Shift keys while selecting the workspace with a mouse click. Then use the Ctrl + Tab keys to switch focus to the next tabbed workspace.

The desktop client and Java Web Start client use these same features to open a workspace, but it is always opened in a new window.

Workspace properties

Every Tivoli Enterprise Portal workspace has properties that control the access and method of display when it is opened. The workspace will open in a new tab on browsers that are set to use tabbed pages when ☑ **Always open workspace in new window** is enabled.

Link target

The link wizard Target Workspace page has an option to ☑ **Always open target workspace in new window**. It shows in the link wizard Parameters page as ▯ *openTargetInNewWindow*. When enabled, this option opens the targeted workspace in a new tab if you are logged on from a tab-enabled browser.

Navigator item find

🔍 **Find** in the Navigator toolbar enables you to locate any Navigator item using simple or advanced search criteria. The Find window has an option to ☑ **Open workspace in new window** that will open the default workspace for the found Navigator item in a new tab if you are logged on from a tab-enabled browser.

On demand

You can open a workspace in a new tab from any context:

Open a workspace
- Ctrl + Shift + click a Navigator item to open its default workspace.
- Click ⊞ **Workspace gallery** and Ctrl + Shift + click the workspace.
- Right-click the active Navigator item, point to **Workspace**, and Ctrl + Shift + click the workspace that you want to open.
- View > Workspace > Ctrl + Shift + click the workspace that you want to open.

Link to a workspace
- Ctrl + Shift + click the link anchor. If there are multiple choices, click the one you want.
- Right-click the source of a defined link (Navigator item, table view row, pie chart slice, bar chart bar, plot chart point, area chart point), point to **Link to**, and Ctrl + Shift + click the link name.

Open a found Navigator item at its default workspace
1. In the Navigator view toolbar, click ⚲ **Find**, enter the find criteria and click **Find**.
2. From the list of Navigator items that is displayed in the **Find results** area, point to a row and Ctrl + Shift + click to open the default workspace for that Navigator item in a new tab.

Open a Navigator view at its default workspace
- In the Navigator view toolbar, click the View ▾ list box and Ctrl + Shift + click the Navigator view.
- View > Navigator View > Ctrl + Shift + click the Navigator view.
- Ctrl + Shift + click a Navigator tab.

Troubleshooting

If the workspace or Navigator view opens in a new window rather than a new tab, review the tab options in your browser to ensure that the tab feature is enabled. If you are using Microsoft Internet Explorer 7, be aware that tabbed workspaces are treated as pop-ups: In the Tabbed Browsing Settings window (Tools > Internet Options > General > Tabs > Settings), ◉ **Always open pop-ups in a new tab** must be selected.

Refreshing a workspace

You can refresh the data that is displayed in the workspace on demand or at a set interval.

About this task

The Tivoli Enterprise Portal client receives monitoring data from monitoring agents whenever you open a workspace that includes query-based views. The default setting for most predefined workspaces is *On Demand*, which means retrieved data remains static until you refresh manually.

Procedure

- To refresh a workspace manually, click ⟳ **Refresh Now**.
- To set a refresh interval, click **View > Refresh Every**, and select one of the intervals: ⏲ 30 seconds; ⏲ 60 seconds; ⏱ 5 minutes; ⏱ 15 minutes; ⏱ 60 minutes; or On Demand .

What to do next

Note: In order for your refresh settings to persist, you must save the workspace and optionally select "Assign as default for this Navigator item" in the **Save Workspace As** dialog.

If the workspace includes any plot charts or area charts, you can refresh those views independent of the rest of the workspace by editing the Style properties of the plotted area to specify a refresh rate.

Be aware that the more frequent the automatic refresh, the more network traffic you create. These requests travel from the portal client to the portal server and to the hub monitoring server before reaching the monitoring agent. They might also pass through a remote monitoring server to reach the monitoring agent. The information is returned by the same route.

Related tasks:

"Suspending and stopping refresh"
If the workspace is set to refresh automatically at timed intervals or it includes event status views, you can suspend refreshes to keep the data from changing while you investigate a problem.

Suspending and stopping refresh

If the workspace is set to refresh automatically at timed intervals or it includes event status views, you can suspend refreshes to keep the data from changing while you investigate a problem.

About this task

When you open a workspace that includes table or chart views, the Tivoli Enterprise Portal receives the most recently sampled monitoring data from the agents. Take one of these steps to suspend data refreshes or to stop receiving the data that populates the workspace.

Procedure

- To suspend automatic refreshing of the workspace, click ▐▐ **Pause Refresh**; click ▶ **Resume Refresh** to turn on automatic refresh again.
- To stop loading the workspace, click ⬤ **Stop**.

Related tasks:

"Refreshing a workspace" on page 41
You can refresh the data that is displayed in the workspace on demand or at a set interval.

Linking to a workspace

Use these steps to link to a workspace that has been targeted from the current Navigator item or view.

About this task

Many monitoring agent products have workspace links available through their predefined workspaces. You can also create and use links to workspaces that follow a logical progression of investigation into performance and operation issues.

Procedure

1. Open the workspace from where you will launch the link. This is the source workspace.
2. Do one of the following, depending on where the link originates:
   - Right-click the current (highlighted) Navigator item.
   - Click the 🖉 link indicator on a table row or the graphic view, then skip to step 4. (A dimmed link indicator means the link is not available from that row.)
   - Right-click a pie chart slice, bar chart bar, plot chart point, table row, graphic view icon, or TMS Infrastructure object.
3. Click **Link To** and click the target workspace in the 🖉 list. The target filter or link expression is used to select the information displayed in the views of the target workspace. If, instead of the workspace opening, you get a "Target not found" message, the definition of the target workspace could not be resolved.
4. If there is more than one workspace you can link to, the Select Target window opens: Select the Navigator item for the workspace and click **OK**.
5. If a message asks you to select a leaf node, click **OK**, then select an item deeper in the tree hierarchy.

Results

The target workspace is displayed. If the link was defined to open the workspace in a new window, it is opened in its own window. If you are using the browser client and your browser supports tabs, the workspace is in a new tab next to the source workspace.

Note: Navigation using ⬅ and ➡ to visited workspaces retains the link context. As an example, consider a link to a workspace from a table row. The row from which you linked is remembered when you revisit the target workspace.

Related concepts:

"Defining links between workspaces" on page 75
The Tivoli Enterprise Portal enables you to define a link from one workspace to another. Then you can quickly jump to a related workspace or a sequence of workspaces that you like to frequently review.

View title bar and toolbar

Every workspace view and the Navigator view has a title bar with some or all of these controls.

✎	Opens the Properties editor to the properties for the view.
⬇ ⬆	Shows or hides the view toolbar. This button does not display if the view has no toolbar.
⊟	Splits the Navigator horizontally to create a new workspace view.
▢	Maximizes a view for a closer look. Click **Restore** to return to the original size. You can save the workspace with the view maximized.
❌	Removes a view. There is no undo for this action except to open a different workspace and answer No when a message asks if you want to save the workspace; -OR- Select **File > Save As** to keep the original workspace with the view intact and create a new workspace without this view.

Most view types have a toolbar for performing specific actions in the view. A common tool is 🔍 **Find**, which is available for finding values in the browser view, notepad view, table view, message log view, and the event console views. Another tool, 🕒 **Time span**, is for specifying the time period to be displayed in a query-based view when historical data is being collected for it.

Setting a time span to display

You can have the Tivoli Enterprise Portal log data samplings into history files or a database for display in a table or chart. These historical data enabled views have a tool for setting a range of previous data samples to be reported.

Before you begin

Historical data collection must be configured and distributed to the managed systems that you are querying data from. Otherwise, 🕒 **Time Span** is not displayed in the toolbar of the query-based view.

Some attributes groups, such as Situation Status and the Windows Event Log, are historical in nature and show all their entries without you having to specify a time span. For these types of attribute groups, you do not need to configure history collection unless you want to roll off to a data warehouse for long-term storage or limit the amount of data displayed. Thus, the time span feature, rather than showing more data, limits the data reported to the time period indicated.

Even if data collection has been started, you cannot use the time span feature if the query for the chart or table includes any column functions. If this is the case, you can select or create another query to enable 🕒 **Time Span**.

About this task

Take these steps to broaden the time range of data beyond the current data samplings.

Procedure

1. Open the workspace containing the chart, table, or relational table-based topology view where you want to see historical data.
2. Click 🕒 **Time Span** in the view's toolbar.
3. Select a time frame: **Real time plus Last _ Hours** (enabled for bar, plot, and area charts only), **Last _ Hours** (or Days, Weeks, or Months, if the data is warehoused), or **Custom**.
4. If you selected **Last** or **Custom**, specify the range of data:
   - Detailed data is all the data collected for the agent.
   - Summarized data is data that is aggregated across the specified time frame. Set the time frame interval, work shift, and the days. This option is disabled if no data warehouse and summarization schedule were configured for this attribute group.
   - To set a custom time frame, click inside the **Start Time** and **End Time** fields to open the date editor. Use the 🔢 spin boxes to adjust the time, year, or month; and click the calendar day. **HH:MM:SS AM/PM** is initially set to the current time.
   - To apply the time span to all views that use the same query as this view, select ☐ **Apply to all views associated with this view's query**. When this

option is enabled, the query is modified to include the time span set here, so any other views using this query report the same time range.

- The **Timestamp** column that is added to the historical view can show either the global timestamp (the default) or local. Select ☐ **Use Hub time** to reflect the time at the Hub Tivoli Enterprise Monitoring Server rather than at the Tivoli Enterprise Portal Server or portal client.

Results

After you click **OK**, the view shows data from the time span you specified. If the view is a table, a timestamp is displayed as the first column and is accurate to the nearest minute; seconds are displayed as 00.

If you see `null` as the value of a table cell or chart point, it means that no value was stored in the Tivoli Data Warehouse. This happens when values that were identified as invalid are reported from a monitoring agent for a given summarization period.

What to do next

The sort function is incompatible with the historical reporting feature. If you are using 🗓 **Time Span** to retrieve historical data, the chart or table is not sorted even if you have specified a sort order in the query. You can still sort a table by clicking a column heading.

After support for an updated product has been applied to the portal server, it is possible to get a request error message about a missing or unknown column name in the view's status bar after you set a time span with ☑ **Use summarized data** selected. Wait until after the next scheduled summarization and pruning procedure has taken place before viewing the summarized data. If need be, you can reschedule summarization and pruning to run sooner.

Related concepts:

"Overview of historical data collection and reporting" on page 217
Use historical data collection and reporting to gather useful metrics about your managed network. You can also use historical data with the chart baselining tools for predictive analysis and in situation modeling for key performance indicators.

Related tasks:

"Discontinuing historical reporting" on page 220
For views that have a time span set to show data samplings over a range of time, you can remove the time span and display real-time data only.

Moving a view

Move a view to a different position in the workspace to take best advantage of the available space.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

About this task

You can easily rearrange the views in a workspace to take advantage of the greater width available below the Navigator view and the narrow area to the right.

Procedure

1. Move the mouse pointer over the title bar of the view you want to relocate.
2. Click and drag the view over the space you want it to occupy. While dragging the title bar, you should see a semi-transparent copy of the view.
3. Release the mouse button to swap the views.

Reordering columns and rows

Change how a table view is sorted temporarily to focus on specific column details or in a more permanent way if it should always appear sorted in this manner. You can also reposition columns for a more logical organization or to emphasize columns of interest.

About this task

You can set the sort order in three ways. Take one of the following steps

Procedure

- Reorder the table view columns or chart view data series by clicking the column heading and dragging to the new position:
  - If you open a workspace and drag the column heading in a workspace table view, the change is temporary unless you save the workspace. You can also right-click a column heading and click 🔒 **Lock this Column** to fix the column and any columns to the left of it so that they do not move off the viewable area of the pane when you scroll horizontally.
  - If you click ✏ **Edit properties** in the table- or chart view title bar and drag a column heading in the ⇶ **Filters** tab, the change is temporary unless you save the workspace.
  - If you click ⒥ **Query Editor** and drag a column heading in the Specification tab, the changes are saved with the query and applied to any views that use that query. Column reordering must be done in a custom query because predefined queries are not editable.
- Sort the row data in a table view:
  - Click ⒥ **Query Editor**, open a custom query, click **Advanced**, and select a **Sort By** attribute from the ▾ list.
  - In a table view, click a column heading to sort all rows by that column in ▴ ascending order, ▾ descending order, and back to the original order the third time you click. You can save the view's sort order with the workspace.

Zooming a chart view

Zoom in on part of a bar chart, plot chart, or area chart to see a subset of the data in greater detail.

About this task

Zooming expands the chart along the category axis; for example, you can zoom a bar chart to show just a few bars, or zoom a plot or area chart to show a shorter span of time. While displaying the chart, complete the following steps to zoom:

Procedure

1. Select an area of the chart by clicking and dragging the mouse pointer. The selected area immediately expands to fill the available view space. Each time you click and drag, the chart zooms to show more detail.
2. Press Esc for each zoom step to undo. The plot area returns to the previous size.

Exporting a query-based view

Save some or all of the data that was retrieved by the query for a view in a text (.txt) file or comma separated value (.csv) file for further manipulation in another software application.

About this task

When you open or refresh a workspace, the queries assigned to the table and chart views are sent to the Tivoli Enterprise Monitoring Server, which returns the requested data from the managed systems. Any filters applied to views reduce what displays, such as only the relevant four columns of a seven-column table or only the processes that are using the most memory.

When you export a query-based view, you can save everything returned by the query before any filters are applied. This includes every column queried or only those you select; every row returned or, if there are multiple pages, only the rows from the displayed page. Also, if you save the data to an existing file, you can choose to append the data or to overwrite the previous version. Take these steps to save the data you select from the view to a .txt or .csv file:

Procedure

1. Open the workspace with the query-based view whose results you want to save.
2. Because the current data display is exported, click ♻ **Refresh** if you want the latest results.
3. If this is a table, pie chart or bar chart view with multiple pages and you want to save a particular page rather than all rows, click ▼ or ▲ at the ends of the vertical scroll bar (or type the number in the *Page* text box) until it displays. A page is 100 rows by default. You can change the view's page size by clicking ✎ 🖳 and changing the **View-level Page Size** value in the Query tab.
4. Right-click inside the query-based view and click 🖳 **Export**.
5. In the Export to Disk window, select the location for the exported data, select a file type, and type a file name.
6. To export the rows that are visible in the view space and all the columns shown in the **Assigned** list, click **OK**; or, to specify otherwise, do any of the following:

   - For the all pages returned by the query, click ◉ **All rows**.
   - To exclude a column from those being exported, select from the **Assigned** list and click ➡ ; to include a column, select it from the Available Columns list and click ◀ . Use Ctrl+click to select multiple columns or Shift+click to select all columns from the first one selected to this one.
   - To change the order, select a column and click ▲ or ▼ .

   If this is a new, unsaved table or chart view, ◉ **Visible rows** is selected for you.

7. If the file you are saving already exists, you can choose to overwrite the file with the current data selection, append the data to the file, or cancel.

Finding text in a browser view or notepad view

Use the find function to locate text in the selected notepad view or browser view.

About this task

Take these steps to find specific text or values in a notepad view or browser view.

Procedure

1. Open the workspace that contains the notepad view to search.
2. Click 🔍 **Find** in the toolbar of the view you want to search. (If the tool is hidden, click ☰ Show toolbar in the view's title bar.)
3. In the Find window that opens, type the beginning or all of the text to find.
4. Select or clear any of the following check box options, then click **Find**.
   - ☑ **Case Sensitive** to find only spellings that match the letter casing as typed. When the check box is disabled, all variations of the letter casing can be found.
   - ☐ **Backwards** to search from the cursor point backwards. When the check box is disabled, the default, the search will be from this point forward.
   - ☑ **Wrap Search**, instead of stopping the search when the end of the document is reached, continues to search from the beginning and stops when it reaches the point where the search began. Clear the check box if you want the search to stop at the end of the document (or the beginning of the document if the search is backwards).
5. Click **Find Again** to search for the next instance.

What to do next

You can click **Close** when you are done or leave the Find window open for this view while you continue working with the workspace.

Finding table data

Use the find function to locate a text or numeric value in the selected table view, message log, or event console view.

About this task

Take these steps to use the Find tool for quickly locating values.

Procedure

1. Open the workspace that contains the view you want to search.
2. Click 🔍 **Find** in the table view toolbar. If the tool is hidden, click ☰ Show toolbar in the view's title bar. The Find window opens with the first two column headings (attributes) filled in and fields for entering up to five values.
3. Click an ▾ arrow and select the column to search within. To clear a column name, select the first entry (which is empty) from the list.
4. Click inside the corresponding **Value** field and enter the value to find.

- The default function is ⓥ **Value of expression** (📅 **Compare Date and Time** for time attributes); the default operator is == . If you need to change the function or operator, click the ⓥ or == button and select from the list.
- You can enter a partial value to find all occurrences of the text or number that begin with those characters. For example, 500 will find 500 and 5000; and West will find West, West006, and Western. You can also use an * asterisk wildcard for any number of characters beginning at that position; use a ? question mark wildcard for a single character at that position. You can also use wildcards with the ⓐ **Scan for string within a string** and ⓐ **Return a subset of the string** functions. See the formula function topics for a description of attribute types and other formula functions, including examples. When using the ? question mark wildcard with the ⓥ **Value of expression** function, append the text with an * asterisk. For example, enter Sa?e* to find Safe and Sale.
- Most special characters are processed correctly. If you attempt to find text that includes a special character and it is not found, but you know it exists in your environment, precede it with the \ backslash escape character.
  These are some of the special characters that require the \ escape character prefix: # pound sign; $ dollar sign; ^ circumflex; \ backslash; (open and) close parentheses; + plus sign; [open square bracket; { open and } close curly braces. For example, \#my_situation to find #my_situation in the Message Log view.
  These special characters do not need the escape character: ! exclamation point; @ at sign; % percentage; & ampersand; = equal sign; < open and > close pointed bracket; _ underscore; - hyphen;] close square bracket; / slash; : colon; and ; semicolon.
5. Select or clear any of the check boxes and click **Find**.
   - ☑ **Case Sensitive** to find only spellings that match the letter casing as typed. When the check box is disabled, all variations of the letter casing can be found.
   - ☐ **Backwards** to search from the cursor point backwards. When the check box is disabled, the default, the search will be from this point forward.
   - ☑ **Wrap Search**, instead of stopping the search when the end of the document is reached, continues to search from the beginning and stops when it reaches the point where the search began. Clear the check box if you want the search to stop at the end of the document (or the beginning of the document if the search is backwards).
6. Click **Find Again** to search for the next instance or enter another search term.

What to do next

You can click **Close** when you are done or leave the Find window open for this view while the workspace is open.

Related concepts:

"Table find examples"

Related reference:

Cell functions
Use the cell formula functions to add filters to queries and query-based views, and to set thresholds for situations and query-based views. The functions that are available depend on the type of attribute (numeric, text, or time) and in which editor you are working: Query, Situation, or Properties (Filters and Thresholds tabs).

Table find examples

The criteria in this example will find cells whose AvgDiskBytes/Write value is at or above 5,000 and that were written on first day of the month (**Return a subset of the string** in the formula is set to start at the fourth character in the time stamp, so it will find 06/01/09 and 10/01/09 but not 06/13/09).

Column	Value
AvgDiskBytes/Write	>= 5,000
Timestamp	== 4,01

The criteria in this example were written for the table in the EIB Change Log workspace to find entries that have been updated and that have "Disk" somewhere in the name. (**Return a subset of the string** and **Scan for string within a string** are described in STR and SCAN.)

Column	Value
Table Name	== IDDefSituation
Operation	== Update
Object Name	SCAN== 'Disk'

Most special characters are processed correctly in the find function. If you attempt to find text that includes a special character and it is not found, but you know it exists in your environment, precede the character with the \ backslash escape character. In this example, the user is searching through Temporary Storage Queues for a Queue ID of #ETC. This is how it is entered in the Value cell:

Column	Value
QueueID	== '\#ETC'
CICS RegionName	

Related tasks:

"Finding table data" on page 48
Use the find function to locate a text or numeric value in the selected table view, message log, or event console view.

Related reference:

Cell functions
Use the cell formula functions to add filters to queries and query-based views, and to set thresholds for situations and query-based views. The functions that are available depend on the type of attribute (numeric, text, or time) and in which editor you are working: Query, Situation, or Properties (Filters and Thresholds tabs).

Printing a view or workspace

You can print the workspace or an individual view and display a preview beforehand.

About this task

Do any of the following to see a print preview or print the workspace or selected view:

Procedure

- To see a print preview of the workspace, click **File** > 🖳 **Print Preview**.
- To print the current workspace, click **File** > 🖨 **Print**.
- To see a print preview of the view, right-click in the view and click 🖳 **Print preview**,
- To print the view, right-click in the view and click 🖨 **Print**.

Results

When you print a table that has more columns than will fit the width of the page, they will wrap to the next page. If there are more rows than will fit on the page, they are displayed on subsequent pages.

What to do next

The printing mechanism used by Tivoli Enterprise Portal is provided by Java, and has some limitations. When you print a notepad view, the text prints as an image. If you want a better quality, try selecting and copying the text (Ctrl+C), then pasting into your favorite text editor for formatting and printing. When you print a workspace or chart, try sizing the application window or view pane smaller to improve the clarity of the printed copy. You can also use the **Print Screen** key.

As an alternative to printing in Tivoli Enterprise Portal, you can export the view to a comma separated value (.csv) file or text (.txt) file, then import it into the program of your choice, and print there.

Related tasks:

"Exporting a query-based view" on page 47
Save some or all of the data that was retrieved by the query for a view in a text (.txt) file or comma separated value (.csv) file for further manipulation in another software application.

Chapter 4. Responding to events

Event indicators such as ⚠ Minor or ✅ Harmless appear on Navigator items when an event is opened for a true situation.

A situation is a test for certain conditions on managed systems that, when met, change the situation status to true. The situation can include a take action command (reflex automation) that will be carried out when the situation is true.

Understanding situation events

The products that run in the Tivoli Monitoring environment come with their own set of situations. These predefined situations serve as models for defining custom situations for your environment.

Situation formula

Situations formulas are constructed of one or more expressions. For example, a situation that checks for free disk space below 20% on a Windows system has an expression that uses the Logical Disk attribute "Free Megabytes" and reads as `Free Megabytes < 20`.The situation will read data samples taken at the managed system at set intervals, such as once a day for the disk space situation in our example.

Other situations might be more elaborate, such as the predefined Windows OS situation, Bottleneck Memory. It embeds two situations: one that tests for excessive memory paging activity (>100 pages per second), and one that tests for processor time over 70%. If both of these situations are true at the same time, the Bottleneck Memory situation becomes true and opens an event.

Situation event indicators

When a situation is associated with a managed system, it also has a state setting that determines which light, is displayed on the Navigator item when the situation becomes true. From highest severity to lowest, they are ⊗ Fatal, ⊗ Critical, ⚠ Minor, ⚠ Warning, ▣ Harmless, ▫ Informational, and ⊕ Unknown. As you move up the Navigator hierarchy, multiple events are consolidated to show only the indicator of the highest severity. Go to the lowest level of the hierarchy in the Navigator and you see the event indicator over the attribute category for which it was written.The situation event console and graphic view also show situation event indicators and enable you to respond to events. The Enterprise Status workspace includes the situation event console view.

Enterprise Status workspace

The initial default workspace is the Enterprise Status workspace. The views in this workspace give an excellent overview of situation events in your monitored environment and their status. One of the views, Situation Event Console, lists the open events and their severities.

Event flyover list

In the Navigator view or a graphic view, you can hover the mouse over an event indicator to open a flyover listing of open situation events with this information:

Event state

Situation name

Name of the system on which the event occurred

Two-letter code for the monitoring agent

Event timestamp

Display item, if one was specified

Situation status icon if: ☑ the event has been acknowledged; ⬤ the acknowledgement has expired; ↻ the acknowledgement was removed before it had expired and the situation is still true; ⚠ the situation is not functioning properly and you will not be able to open its event results workspace; or ⬤ the situation has been stopped and you will not be able to open its event results workspace.

The ✎ link is to the event results workspace for the situation.

The pop-up menu for the event items has options for working with the situation and event:

- ✎ Edit the situation
- ▶ Start the situation
- ⬤ Stop the situation
- ⚡ Quick acknowledge the event
- ☑ Acknowledge the event
- ☑ Close the event (pure events only)

Event results workspace

The event results workspace shows the values of the attributes when the situation fired and their current values. It shows any expert advice the situation author might have written and any hypertext links to go to for more information. The take action view enables you to select or define a command to be invoked at the managed system.The table views showing the initial situation values and their current values also have pop-up menus that include the option to launch an application. If you have a favorite diagnostic tool, you can start it from here.

Event acknowledgement

When multiple users are monitoring the network for events, one of them can acknowledge an event to indicate it has been seen and the problem is being worked on. Acknowledging an event changes its event status from Open to Acknowledged until the acknowledgement expires or until the situation is no longer true.You can add notes to an acknowledgement and attach files that are pertinent to the event.

Automated responses

Automated responses to events can improve the efficiency of systems management by reducing user workload and errors and allowing quicker responses to problems. Tivoli Enterprise Portal provides two levels of automation: *reflex* and *advanced*.

- Reflex automation provides simple automated system actions. It allows you to monitor a condition on a particular system and to specify a command to execute there. The command can be a single action or a script of commands. Tivoli Enterprise Portal receives no feedback after it sends the command or notifies the user.

  Reflex automation is implemented by adding an action command to a situation definition that runs when an event is opened.

- Advanced automation allows you to implement more complex workflow strategies than you can create through simple reflex automation. *Policies* are used to perform actions, schedule work, and automate manual tasks. A policy comprises a series of automated steps called *activities*, which are connected and can branch to create a *workflow*. After an activity is completed, a return code is received with feedback and the next activity prescribed by the feedback is carried out.

  With policies you can monitor multiple conditions simultaneously on any number of systems and have selected activities take place when specific conditions occur. Many monitoring agents come with predefined policies that you can use as is or modify for your environment.

  You can also create your own policies. The Workflows editor is provided for designing and managing policies.

Opening the situation event results workspace

When the conditions of a situation have been met, event indicators appear in the Navigator and in the graphic view. You can investigate the event by opening its workspace.

Procedure

- Open the event results workspace from the event flyover list:
  1. In the Navigator or graphic view, move the mouse over the event indicator (⊚ Fatal, ⊗ Critical, ⚠ Minor, ⚠ Warning, ▫ Harmless, ▫ Informational, or ⊚ Unknown) to open a flyover listing of the situation that caused the event. The Navigator shows one event indicator for an item. If more than one situation is true, the event indicator with the highest severity is displayed. Entries in the event flyover list are grouped by severity. Each event reports the situation name, the computer where the event occurred, the two-letter code of the monitoring agent (such as LX for the Linux OS agent), the event timestamp, and the display item if one was specified. As well, if any of the following changes have occurred, a situation status icon displays at the far right:
     ☑ The event has been acknowledged.
     ⊕ The acknowledgement has expired.
     ↻ The acknowledgement was removed before it had expired and the situation is still true.
     ⚠ The situation is not functioning properly and you will not be able to open its event results workspace.
     ⬤ The situation has been stopped and you will not be able to open its event results workspace.
  2. In the flyover list, click ✎ next to the situation to open its event results workspace.
- Open the event results workspace from the situation event console view:
  1. Right-click the row of the event.
  2. Click **Situation Event Results**.

Results

A new Navigator item is displayed below the one associated with the event and the event results workspace opens. The event results workspace contains four views to help you investigate the condition and determine what action to take next:

⊞ Initial Situation Values

> The values are from the data sample that triggered the event. Move the mouse pointer over a highlighted value to see the situation formula. The value shows as "Unknown" and the formula does not display if the situation expression uses a group function: CHANGE, PCTCHANGE, AVG, COUNT, MIN, MAX or SUM.

⊞ Current Situation Values

> The values are from a new data sample that was taken after you selected the event results workspace. If the situation is still true, you will see the same cell highlighting as is shown in the Initial Situation Values view. If this is a pure event, you will see no data. This is because pure events are unsolicited notifications. Situations for pure events have no sampling interval, thus there is no constant metric that can be monitored for current values.
>
> If the expression that is responsible for a situation event uses a group function, the situation event results workspace shows neither the attribute value threshold color nor the expression that should display when you hover over an initial situation value.

⚡ Command View

> Use the command view to send a take action command to an application started on the managed system of your choosing.

ⓘ Expert Advice

> If the situation author wrote any expert advice, it is displayed here. This view opens to a web page if the advice is a URL reference.

What to do next

Compare the current situation values with the initial situation values that caused the event and read the advice for any explanation and possible solutions.

You can right-click the ▦ situation workspace Navigator item to open the pop-up menu. The menu includes functions for responding to the event:

✎ Edit Situation

> Opens the Situation editor so you can see and edit the situation definition, distribution, expert advice, take action command, or until modifier.

☑ Acknowledge

> Opens the Acknowledgement window so you can acknowledge that you are taking care of the event.

☑ Close Situation Event

> If this is a pure event, you will see this option to close the event and remove the event indicator. If the condition goes from true to false and back to true again the next time the situation is evaluated, the indicator will reappear. If the condition remains true, the indicator is not displayed; you must restart the situation if you want to see the indicator.
> If you close an event manually, you cannot acknowledge it.

☒ Remove Event Item

> Removes the event item from the Navigator. To see it again, you must select the true situation from the event flyover list.

User ID permissions for Events are Acknowledge and Close; and those for Situations include Start or Stop, View, and Modify. If any of these fails to appear in the pop-up menu, your user ID does not have permission for the function.

If you have workspace author mode permission, you can also edit the workspace to add any views that are helpful for diagnosing an event for this particular situation. For example, if historical data is collected for an attribute group, you might want to create a plot chart and add a historical baseline to help to get a broader view of past and current values for trend analysis.

Enterprise Status workspace

At the highest level of the Navigator Physical view is the ▦ Enterprise item. The default workspace for the Enterprise is called Enterprise Status.

The Enterprise Status workspace contains four views:

Situation Event Console

The situation event console inserts a row when the status of a situation changes: Open (including expired and reopened events), Problem, and Acknowledged. The Owner column shows who acknowledged the event, and you can acknowledge an event from the pop-up menu. Link indicators in this view and in the My Acknowledged Events view open the "Event Details workspace" on page 58 or you can right-click a 🖉 link indicator and select another of the linked workspaces.

My Acknowledged Events

The My Acknowledged Events table view displays the events assigned to the current user. This view shows both open and closed events. Closed events are displayed to enable you to access the notes and actions taken against them.

For each event record listed, you can click the 🖉 link indicator at the beginning of the event row to open the "Event Details workspace" on page 58 with similar events based on the situation name; or you can right-click the link indicator to see the other Event Details workspaces. The Event Details workspace shows any notes and attachments for the acknowledged event, along with other related events and a link to the IBM Integrated Service Management Library (http://www.ibm.com/software/brandcatalog/ismlibrary) where you can find analytical tools for situations.

If the event is from an incident management agent (such as IBM Information Management), you will also see a Reference ID column and, when you open the pop-up menu to select a linked workspace, a link to the Incident Details workspace.

Message Log

The message log view displays a row for every situation in the monitored network whose status changes to any of the following: *Open*, *Closed*, *Stopped*, *Problem*, *Acknowledged*, *Expired*, or *Reopened*.

Open Situations Counts - Last 24 Hours

The Open Situations Counts - Last 24 Hours is a chart view which displays a bar for every situation in the monitored network that has become true in the past 24 hours. The bar size indicates the number of times a situation event has been opened during this time.

More predefined workspaces

More predefined workspaces for the Enterprise Navigator item are accessible through the ⊞ **Workspace Gallery**, or by selecting the **Workspaces** menu option. In addition, your monitoring product might have workspaces at the Enterprise level. Consult your product help for more information.

Event Details workspace

The Enterprise Status workspace (default for the Enterprise Navigator item) includes the Situation Event Console and My Acknowledged Events views, each of which has links to the Event Details workspace. The Event Details workspace provides detailed information about the event row and has four views.

Selected Event

The Selected Event view provides reference information about the event. You can click the 🖉 link indicator to see the other two Event Details workspaces, then select one to open it.

Similar Events by...

This is the view that distinguishes each type of Event Details workspace. The events will be similar by situation name, by source, or by resource. This view enables you to identify event patterns to help diagnose the root cause, whether by condition, by managed system or by the specific resource being monitored.

Similar Events by Situation Name
This view lists all the events for the same situation. Such events would occur when the situation is true on more than one managed system or the situation is true on the same managed system more than once (a display item was defined in the Advanced Situation options window).

Similar Events by Source
This view list all the events from the same managed system as the event you are viewing.

Similar Events by Resource
This view displays all the events for situations that were defined with the same attribute group as the one used for this situation.

Event Notes

The Event Notesdisplays any notes and files that have been applied to the event acknowledgement. You can access any attached file by clicking the Attachment hypertext link. The file is launched in your default browser or in the default viewer configured for the file's MIME type. You can print the notes for an event by clicking 🖶 **Print**. If you want to update the notes for an event, right-click the event in the Selected Event view and click 📄 **Event Notes**. Click 🔄 to view the updates.

Event Tools

The Event Tools browser view provides links to additional resources for managing events. There are downloadable tools, some of which can be used to troubleshoot any problems you have with events. For example, this view provides a link to the Log and Trace Analyzer, which you can use as part of an autonomic event management solution. This view also provides a link to the IBM Integrated Service

Management Library (http://www.ibm.com/software/brandcatalog/ismlibrary) website, from where you can access additional problem determination information and tools.

Acknowledging an event

When an event is opened for a true situation, you can acknowledge it. This notifies other users that you have taken ownership of the problem related to the event and are working on it. The acknowledgement window records the time of the acknowledgement and the ID of the user.

Before you begin

Your user ID must have Acknowledge permission for Events to use this function.

About this task

Acknowledgements enable users who are monitoring events to communicate their working status. Follow these steps to acknowledge an open event or one whose acknowledgement has expired or been canceled before the event has been closed.

Procedure

1. Do one of the following:
   - Move the mouse over an event indicator (such as ⊗ , ⚠ , or ▫ in the Navigator to open the flyover list of open events, then right-click the event you want to acknowledge.
   - If you have opened the event results workspace, right-click the 🗹 event item in the Navigator.
   - In the event console view, right-click a row.
2. Point to **Acknowledge Event** and click either ☑ **Quick Acknowledge**, to acknowledge the event with the default settings, or ☑ **Acknowledge** to open the Acknowledge window. A quick acknowledgement creates an acknowledgement with no expiration and inserts a time stamp, your user ID, and the statement, "Event quick acknowledged, ownership assigned" in the Notes field.
3. If you selected ☑ **Acknowledge** in the previous step, do any of the following:
   - Adjust the **Expiration** settings for the time when the acknowledgement should expire. If you select ◉ **Expire at a specific time**, click the ▾ list box, adjust the year, month, or time, then click the expiration date to close the calendar. Select ☑ **Use Server time** to use the time zone of the Tivoli Enterprise Monitoring Server to which the monitoring agent is connected.
   - Write notes about the event in the **Add Notes** text box. After you click **Apply** or **OK**, the text will be moved to the **Notes** area and preceded with a time stamp, your user ID, and the statement, "Event quick acknowledged, ownership assigned".
   - Attach a file that is applicable to the event by clicking 📎 **Add an attachment** and selecting the file.
   - If your environment includes an incident management agent, such as IBM Information Management, and the event corresponds to an incident, enter the incident identifier in the **Reference ID** field, up to 32 characters.
4. When you are finished, save your notes and settings by clicking **Apply** to keep the Acknowledgement window open or **OK** to close the window.

Results

Any notes just written and attachments added are displayed in the **Notes** area, including the date and time each note was added, the user ID of the author, and hypertext links to the attachments.

What to do next

The acknowledged status of the event is indicated by the ☑ Acknowledgement icon in the event flyover list. The message log, situation event console, and common event console views also show "Acknowledged" in the Status column. When the event is closed, the acknowledgement is removed. If you want to review the acknowledgement after it has been closed, you can access it through the view of the Enterprise Status workspace. If the acknowledgement expires before the event has been closed, the event status changes to ☺ Reopened.

Acknowledgement status

After you acknowledge an event, its status changes to ☑ Acknowledged and is displayed in several places.

- Navigator item for the event (which is displayed when you open the Situation Event Results workspace)
- Situation flyover list at the end of the row
- Status column of the message log view, situation event console view, and in the Enterprise Status workspace view named My Acknowledged Events.

The Acknowledged status changes to Closed when the situation is no longer true.

If the situation is still true when the acknowledgement expires, the status will change to ● Expired. If the situation is still true and you cancel your ownership (remove acknowledgement) the status will change to ☺ Reopened.

Although you can acknowledge a pure event, the acknowledgement does not stop the arrival of more pure events.

🏢 Enterprise
   🖥 UNIX Systems
      📁 BigComputer
         ◎ Oracle
            🖥 Cache Totals
               ☑ Oracle_Library_Reloads_Pct_Crit
            ⇟ More...
        ⇟ More...
     ⇟ More...
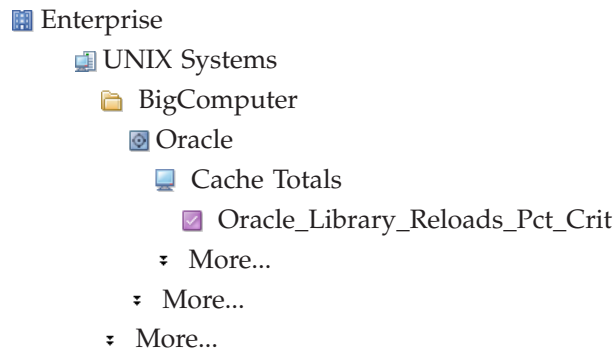
Reviewing the acknowledgement and adding notes
About this task

The Acknowledgement feature includes a window for adding notes about the event and for attaching files. You can review all the notes that have been written during the lifetime of an open event, including notes that were added before the acknowledgement expired and those that are added automatically when an acknowledgement is removed. And you can append notes and attachments at any time.

Procedure

1. Do one of the following:
   - Move the mouse pointer over an event indicator (such as ⊗ or ▣) in the Navigator view to open the flyover list of open events, then right-click the acknowledged event.
   - If you have opened the event results workspace, right-click the ☑ event item in the Navigator.
   - In the situation event console view, right-click a row where the Status shows as Acknowledged.
   - In the Enterprise Status workspace, right-click a row in the My Acknowledged Events view.
2. Click 📋 **Event Notes**

Results

The Event Notes window opens for you to review and add notes for this event. You can also attach files.

What to do next

You can open the window to review the notes and open attachments for the event at any time, event after the event has been closed.

Adding notes to an acknowledged event

About this task

You can add notes when you acknowledge an event or after an acknowledgement has been created (but not to events that are in the Open or Closed state) by you or another user. **Notes** enable you to communicate what steps you have taken thus far to resolve the event or to provide details on the steps you believe must be taken.

Procedure

1. Do one of the following:
   - Move the mouse pointer over an event indicator (such as ⊗ or ▣) in the Navigator view to open the flyover list of open events, then right-click the acknowledged event.
   - If you have opened the event results workspace, right-click the ☑ event item in the Navigator.
   - In the situation event console view, right-click a row where the Status shows as Acknowledged.
   - In the Enterprise Status workspace, right-click a row in the My Acknowledged Events view.
2. Click 📋 **Event Notes**. The **Event Notes** window opens for you to review and add notes for this event. You can also attach files.

Results

You can continue to open the window to review the notes and open attachments for the event at any time, event after the event has been closed. This is done through the pop-up menu of the view.

Removing an acknowledgement

About this task

If you acknowledged a situation event and must release your ownership before the acknowledgement has expired or the event has closed, you can remove the acknowledgement from the event.

Procedure

1. Do one of the following:
   - Move the mouse over an event indicator (such as ⚠ or ▣) in the Navigator view to open the flyover list of open events, then right-click the acknowledged event.
   - If you have opened the event results workspace, right-click the acknowledged ☑ event item in the Navigator.
   - In the situation event console view, right-click a row where the status is Acknowledged.
   - In the Enterprise Status workspace, right-click a row in the My Acknowledged Events view.
2. Click **Acknowledge Event** > 🧩 **Remove Acknowledgement**. Removing the acknowledgement changes the event status to Reopened and removes the user ID from the Owner column.

   A time stamp and user information is attached to the event as a note, along with an audit note signifying that the acknowledgement was removed.

Closing a situation event

When the condition of a situation is met, the situation becomes true and an event indicator is displayed over the system name and all related items in the Navigator. For pure events, you can turn off the indicator by closing the event. The situation continues to run but the indicator disappears until another pure event takes place.

Before you begin

In closing an event, you also have the opportunity to add notes and attachments to explain what action you took to resolve it. This can be helpful should a similar event occur in the future. If the event was not previously assigned to an owner, the user who closes it is assigned as the owner.

To use this function, your user ID must have Close permission for Events.

About this task

Pure events are indicated in the **Type** column of the message log view and situation event console view. As well as closing pure events, you can also close events of sampled attributes if the situation has been stopped or deleted. Close an event from any of the following locations:

Procedure

- **Situation flyover list:** From the event flyover list, do the following to close a pure event or one whose situation has been stopped or deleted:
  1. In the Navigator, move the mouse pointer over the event indicator to open a flyover listing of the true situations.

2. In the flyover list, right-click the situation whose pure event you want to close or the situation that has been stopped or deleted and whose event you want to close.

3. Click **Close Event** → ☑ **Close Situation Event** or **Close Event** → **Close Situation Event with Notes**.

- **Event Navigator item:** From the Navigator item for an event, do the following to close a pure event:

   1. In the Navigator, move the mouse over the event indicator to open a flyover listing of the true situations.

   2. In the flyover list, click a situation name to open its event results workspace. The Navigator focus moves to the event location and shows the situation name, and its event results workspace opens.

   3. In the Navigator, right-click the ▨ event item and click **Close Event** → ☑ **Close Situation Event** or **Close Event** → **Close Situation Event with Notes**.

- **Situation event console view:** From the situation event console view, do the following to close a pure event or one whose situation has been 🔴 stopped or deleted:

   1. In the workspace where the situation event console view displays, right-click the pure event (or the situation that has been stopped or deleted) you want to close. Pure events are indicated by ⊞ Expand in the first cell and *Pure* in the **Type** cell of the event row.

   2. Click one of the following options:

      – **Close Event** → ☑ **Close Situation Event**

      – **Close Event** → **Close Situation Event with Notes**

Results

The indicator disappears until another pure event arrives.

What to do next

Here are some notes about what you can and cannot do after an event has been closed:

- If you close an event, you cannot acknowledge it.
- You can open the notes and any attachments for a closed event through the pop-up menu of the view.
- If you know you need to take an action that will trigger an event, such as stopping an application, but you do not want notification, you can take preemptive measures to stop an event from opening: You can stop the situation then restart it when you want it to begin running again; or, if you can schedule the time when you don't want the situation to run, you can add a time expression to the situation.

Turning off the sound

When an event is opened for a situation that has a WAV file assigned, the sound plays while the event is open.

Procedure

- To turn the sound off, click ◁ **Turn Sound OFF**.
- To turn the sound back on, click ◁)) **Turn Sound ON**.

What to do next

Use your Windows Sound and Audio Devices or Multimedia utility, accessible from the Control Panel, to adjust the volume of event sounds.

Event reporting

For situations that are associated with Navigator items, alert indicators are displayed whenever an event is opened. Several workspace views are also available for displaying situation and event status.

Event views present alternatives to event indicators in the Navigator. Message log and situation event console views update automatically to show new events as they arrive and changes in event status as they occur. Graphic views give you a pictorial alternative to the Navigator for indicating alerts.

Related tasks:

"Adding a graphic view" on page 146
The graphic view provides a canvas for you to cover with icons that represent the Navigator items and that overlay the icons with alert indicators when events are opened. You might create, for example, a graphic view showing your company's floor plan or organization chart.

"Adding a message log view" on page 142
The message log gives you an overview of changes in situation status on your monitored network. You can add the view to any workspace.

"Adding a situation event console view" on page 152
You can add a situation event console view to list each alert for situations associated with the current Navigator item and those it subsumes.

"Adding a common event console view" on page 158
Add the common event console view to a workspace to get a single, integrated display of events from multiple event repositories, including Tivoli Monitoring, Tivoli Enterprise Console, and Tivoli Netcool/OMNIbus.

"Adding a Tivoli Enterprise Console event viewer" on page 164
If your environment includes the Tivoli Enterprise Console, you can add it to any workspace to integrate alarms and events from network management operating systems and correlate them with other events.

"Adding a Universal Message Console view" on page 145
Create a Universal Message Console view to see the messages produced at the hub Tivoli Enterprise Monitoring Server.

Chapter 5. Customizing workspaces

Your Tivoli Enterprise Portal product comes with a set of predefined workspaces for every item in the Navigator Physical view. You can use these as a starting point for creating and customizing new workspaces to suit your environment.

Custom workspace options and guidelines

The Tivoli Enterprise Portal has over a dozen types of views that you can populate a workspace with. Over half of the views are query-based and have a number of embellishments that you can use to observe potential and existing threshold conditions and other metrics.

Workspace changes

You can make these types of changes to a workspace:

- Split a view horizontally or vertically into two separate views.
- Change the dimensions of the views by dragging the borders between them.
- Rearrange views by clicking a view's title bar and dragging it another view space. When you release the mouse button, the views switch position.
- Maximize the view and, if you like, save the workspace while it is still full-screen.
- Change a view to a different type, such as from a pie chart to a gauge.
- Edit the workspace 🖼 properties, such as to make it the default workspace.
- Edit the ✎ properties of the individual views to control their styles and, in table and chart views, what data they display.
- Build a link to another workspace so that you can jump to it from this one quickly.
- Save a copy of the current workspace and edit the copy.

Any changes you make to a workspace are available only to you; no one else will see your changes. The exception is with workspaces that were updated by an administrator while in workspace administration mode, where the changes to that workspace are shared with all users who log on to the same portal server.

Predefined workspaces are protected from editing. You must use **File > Save As** to create your own copy for customization.

Design guidelines

These are the changes you can make to a workspace that are saved with it:

- Display a view full screen.
- Resize views by dragging the borders.
- Swap location with another view by dragging by its title bar to the new position.
- Workspace link definitions.
- Arrange the order of workspaces in the workspace gallery. This gets saved with all workspaces for a Navigator item.

Here are some guidelines to help you plan workspaces:

- Tailor the content of a workspace to the Navigator level: Workspaces at the top of the hierarchy can show summary information; those at lower levels can provide more in-depth information.
- Provide summary information at the agent level of the Navigator Physical view.
- Use custom queries that pre-filter the data used in a table or chart to ensure fastest data retrieval and that you see no extraneous data. And use as few different queries as possible in one workspace.
- Design for a specific monitor resolution. If you are creating in workspace administration mode for multiple users, use 1024 x 768, which is the lowest resolution on which the Tivoli Enterprise Portal can run.
- Do not crowd too much information in one workspace. Instead, have multiple workspaces at one level or links from one workspace to the next.
- For table views with many columns, arrange them so the important ones are visible. Consider removing some columns by pre-filtering (Properties - Query) or post-filtering (Properties - Filters).
- Consider combining tables with charts in a workspace for the most effective style layout.

Tips

- If you edit a predefined workspace or one created in administration mode, you can restore the original through **File > Restore Original Workspace**. The customized version is overwritten by the original. If you want to keep both the customized workspace and the original, save the customized workspace with a new name before restoring the original.
- The managed systems available for reporting are those assigned to the Navigator item. If you are not sure which are assigned, select the Navigator item, right-click the item, and click 🔢 **Properties**.
- The CLI has several tacmds specific to workspaces:

  deleteWorkspace to delete a workspace by its objectid from the Tivoli Enterprise Portal Server

  listWorkspaces to list all of the workspaces, including objectid, on the portal server

  exportWorkspaces to export one or more portal workspace definitions to a file

  importWorkspaces to import the workspaces definitions in a file into the portal server

Workspace Author Mode permission

Your user ID requires Workspace Author Mode permission to create and maintain workspaces, including links. If the main toolbar is disabled (except Refresh, Back, Forward and Stop), as well as the split and remove buttons on the view title bar, you do not have this permission.

Tutorial: Defining a workspace

Use the tutorial to get hands on practice and experience defining a workspace.

In the following exercises you will add new views to an undefined workspace, tailor them with the Properties editor, save the workspace, and, finally, edit the workspace properties.

If the split view and view tools are disabled (dimmed), your user ID does not have Workspace Author Mode permission and you will not be able to complete this lesson.

Split a view

1. Open an undefined workspace.

   An undefined workspace is one that has not been provided with your monitored application or defined by you or an administrator. You can find undefined workspaces initially at the system level (computer or image name is displayed) of the Navigator Physical view, and, if you have custom Navigator views, wherever you have created child Navigator items. A workspace that has not been defined yet will have a browser view with links to this lesson and other relevant topics, and it will have a notepad view.

2. In the notepad view, click ▥ **Split Vertically** to split it into two separate views.

   The notepad view splits in half vertically, with a duplicate copy in the new view space. Notice that every view in the workspace, including the Navigator, has ▥ and ▤ tools, as well as ▨ (except the Navigator) to delete the view and expand the adjoining view to fill the space.

Change the view type

Change the new view to a table view:

1. Click ▦ **Table**.
2. Click inside the notepad view.

   As you move inside the workspace the mouse pointer changes to the table view icon (Linux shows a hand icon). After you click, the new table view replaces the notepad view. Because a query has never been associated with this view, no data is displayed in the table. You will need to select a query from the Properties editor.

3. When a message asks you to assign the query now, click **No**.

Now add a message log view to the original notepad view:

1. Click ▤ **Message Log**.
2. Click inside the notepad view.

   The message log replaces the notepad view.

Tip: If you select the wrong view or change your mind about the view you chose, press Escape or click somewhere in the toolbar.

Swap views

You can easily move a view to a new position by dragging it by its title bar to the view space it should occupy. When you release the mouse button, the views swap position.

1. Point to the message log view's title bar.
2. Click and drag it to the table view.

   As you drag the view, it appears semi-transparent.

Edit view properties

Edit the properties for the table view to select what data should display.

1. In the table view title bar, click ✎ **Properties**.

2. In the **Query** tab, ▣ **Click here to assign a query**.

   Queries determine from which attribute group to draw data. The Select Query window opens.

3. In the Query tree, expand the **Tivoli Enterprise Monitoring Server** branch, then **Managed System**.

4. Select ▣ **Current Managed Systems Status** and click **OK**.

   Back in the **Query** tab, you can see a preview of the table.

5. Click the ▣ **Filters** tab and clear the boxes for these columns (attributes): **HostAddress**, **Affinities**, **Host Info**.

   Filters is where you specify the attributes to include in the table and, if you like, show only those values that fall within a range.

6. Click **Test** to see a preview of your changes.

7. Click the ▪ **Thresholds** tab to specify a threshold.

8. In row 1, click the colored column and select ▣ **Informational**.

   This will be the background color of the cell when the threshold expression becomes true.

9. In row 1, click the Status cell, then the ▾ arrow, and select OFFLINE from the list.

   Any managed systems that are offline will have this background color.

10. Click the ▱ **Style** tab to change the appearance of the table.

11. In the **Header**, type the following title in the Title text box: `Tivoli Enterprise Monitoring Agent Status`.

12. Click **OK**.

Save a new workspace

A new workspace can be saved by following the procedure below:

1. Select **Save Workspace As** from the **File** menu.

   If you are running Tivoli Enterprise Portal in browser mode, be sure to select its File menu and not browser's File menu.

2. Type `My Workspace` in the **Name** field.

3. Select **Assign as default for this Navigator Item**.

4. Click **OK**.

5. Take note of which Navigator item `My Workspace` is associated with.

6. Now test your workspace by opening another workspace, then clicking this Navigator item again.

Edit workspace properties

The workspace properties can be edited by completing the following procedure:

1. With My Workspace open, click ▥ Properties.

   When you open the Properties editor for a workspace from the toolbar as opposed to the pop-up menu, the properties for the workspace display.

2. Clear **Assign as default for this Navigator Item**.

3. Click **OK**.

4. Click ▤ **Save**.

5. Test the new setting by opening another workspace, then clicking this Navigator item again.

My Workspace does not display. Instead, you see the original undefined
workspace, which is the default workspace for this Navigator item.

6. To open My Workspace, right-click the Navigator item, click **Workspace**, and
select "My Workspace" from the list.

Opening a new window

Open a new window.

1. In desktop mode, click new windowNew Window.

   The new window is a duplicate of the original. Any previously visited
   workspaces are retained from the parent window; use Go back and Go forward
   to revisit them. Further navigation to other workspaces in either window,
   however, is independent of the other window.

Note that you can close duplicate windows (click File > Close) or the original; the
work session remains active as long as one window is open.

Note: To view two workspaces at the same time, create a new window. Any
changes you make in this or the other window (or windows, if you open more) are
saved with the application.

Tip: You can set the properties of a workspace to open in a new window
automatically whenever it is selected. If you are running the Tivoli Enterprise
Portal in browser mode from a browser that is enabled for tabs, the workspace
opens in a new tabbed page.

Conclusion

In this lesson you learned some of the techniques used to define and customize
workspaces:

- Split a view.
- Change to a different view: table and message log.
- Swapping views by dragging one view over the other.
- Edit the properties for a view: Select a query for the table, then filter it and add
  a threshold and title.
- Save the workspace, making the new workspace the default for the Navigator
  item.
- Edit the properties of the workspace to remove the default status.

Starting workspace administration mode

Typically, workspaces you create are saved with your user ID and no other users
can see them. You can turn on administration mode so that all workspaces you
edit and save while in this mode are shared with all users who log on to the same
Tivoli Enterprise Portal Server.

Before you begin

In workspace administration mode, the changes you make to link definitions,
launch application definitions, and terminal emulator scripts are also shared with
other users. If users edit a shared workspace, the edited workspace becomes their
personal copy. Any changes made to that workspace in administrator mode will no
longer be received by the user. All other users, including the administrator,
continue to use the shared (global) version.

To use this function, your user ID must be SYSADMIN or created from that administrator ID and have View and Administration Mode Eligible permissions for User Administration.

Procedure

1. Click 👤 **Administer Users** to open the profiles for all users and user groups.
2. In the 👤 **Users** list, click your user ID.
3. In the 🔑 **Permissions** tab, scroll down the Authorities tree and click 👤 **Workspace Administration**.
4. Enable ☑ **Workspace Administration Mode**. (Clearing the check box turns off workspace administration mode.)
5. To save your changes, click **Apply** to keep the window open or **OK** to close it.

Results

The title bar and status bar display ***ADMIN MODE*** after the user name to indicate that your user ID is in workspace administration mode. Any changes that you make to workspaces (including the workspace gallery), link definitions, launch application definitions, and terminal emulator scripts while in workspace administration mode will be shared with other users connected to the same portal server as soon as you save your changes. Likewise, these changes cannot be removed except from within workspace administration mode.

In workspace administration mode, you can access all workspace definitions associated with a particular Navigator item including those that are marked as ☑ **Only selectable as the target of a Workspace Link**. This enables you to modify the target workspaces directly without having to first traverse the link. If a target workspace is dependent on information from the source workspace, you will get display errors when you open the workspace through the **Workspace gallery** rather than the **Link To** list.

Private workspaces that were created by a user outside of workspace administration mode cannot be seen while in administration mode, even if you are that user. For example, if you created private workspace "My charts", you will not see "My charts" in the list of available workspaces while you have workspace administration mode turned on.

What to do next

When editing a workspace in administration mode, select ☑ **Do not allow modifications** in the workspace properties. Otherwise, if a user edits that workspace, you no longer control the workspace and cannot override their changes with any changes you make.

When you are finished configuring workspaces for all users, turn off workspace administration mode by repeating these steps and clearing the **Workspace Administration Mode** check box.

Workspace properties

Every workspace has a name and general characteristics defined in the workspace properties.

When you save a new workspace with **File** ‣ **Save Workspace As**, click 🔳
Properties, or, if the Properties editor is already open, click the top item in the tree
to see the workspace properties.

Workspace Identity

Name The title of the workspace. The workspace name is displayed in the
title bar or browser mode banner, in the **View** ‣ **Workspace** list, or
in the **Link To** list menu, depending on how you have set the
workspace options (see below). Do not use an & ampersand or #
pound sign in the title.

Description
Optional. The description can be a full text description of the
workspace.

Workspace Options

☐ **Assign as default for this Navigator Item**
Makes this workspace the default for the Navigator item (the one
that opens when you click the Navigator item).

☐ **Assign as Home Workspace**
Specifies this as the workspace that opens whenever you start
Tivoli Enterprise Portal, much like your browser home page.

☐ **Do not allow modifications**
Protects the workspace from permanent changes and must be
selected when you are editing in Workspace administration mode.
Users who have access to this workspace can change it temporarily,
but if they attempt to 💾 save, the Save Workspace As window
opens so that it is saved as a new workspace with a different
name. Also, when this item is enabled you cannot delete the
workspace.

☐ **Only selectable as the target of a Workspace Link**
Marks the workspace as hidden except when it is a possible Link
To destination. The workspace does not display in the list of
available workspaces for a Navigator item.

☐ **Always open workspace in new window**
When enabled, this option causes the workspace to open in a new
window, or in a new tab if your browser supports tabbed pages.

Creating a new workspace

You can create a new workspace from any existing workspace. The new workspace
is associated with the same Navigator item as the original workspace. 👤 You
cannot save a new workspace if your user ID does not have Workspace Author
Mode permission.

About this task

To create a new workspace:

Procedure

1. Open a workspace at the Navigator level for which you want the new
   workspace to be available.
2. Click **Save Workspace As** in the **File** menu.

3. Type a name for the new workspace. The workspace is saved with the new name and the original left unchanged.
4. Do any of the following to customize the workspace:
   - Add a new view by splitting an existing view into two separate views.
   - Change the dimensions of the views by dragging the borders between them.
   - Change a view to a different type, such as from a notepad view to a browser view.
   - Rearrange views by clicking a view's title bar and dragging it another view space. When you release the mouse button, the views switch position.
   - Edit the general ⊞ Properties of the workspace, such as making it the default workspace.
   - Edit the ✎ Properties of the individual views to control what data to display and the styling.
   - Build a link to another workspace so that you can link to the other workspace from this one at any time.
5. When you are finished editing a workspace, click ▤ **Save**. The saved workspace is associated with your Tivoli Enterprise Portal user ID. Users with a different ID will not see your changes.

   The workspace is available to all systems with items of the same affinity. For example, if you create a new workspace for your Windows agent then all systems with that monitoring agent will be able to open the new workspace.

Saving a workspace

You must save a workspace to maintain the changes you make to the workspace properties; otherwise, the changes you have made are lost.

Before you begin

Your user ID requires Workspace Author Mode permission to create and maintain workspaces, including links. If the main toolbar is disabled (except Refresh, Back, Forward and Stop), as well as the split and remove buttons on the view title bar, you do not have this permission.

About this task

Whatever changes you make to the workspace properties, such as adding a view or defining a link, are only temporary. They will be lost when you open another workspace or exit Tivoli Enterprise Portal unless you save the workspace.

The saved workspace is associated with your Tivoli Enterprise Portal user ID. Users with a different ID will not see your changes unless you were in workspace administration mode when you authored them.

Procedure
- To save the current workspace, click ▤ **Save**.
- To save the current workspace as a new workspace and leave the original unchanged, click **File** > **Save Workspace As**, type a name for the new workspace, check any of the workspace options that you want to apply, and click **OK**. The workspace is saved with the name and options you specify, and is available to all systems with items of the same type. For example, if you create a new workspace for your Windows agent then all systems with that monitoring agent will be able to open the new workspace.

Editing a workspace

Use the Properties editor to change the general characteristics of a workspace and to edit the style and content of any of its views. You can delete or restore the original version of an edited workspace through the File menu.

Before you begin

To edit a workspace, your user ID must have Workspace Author Mode permission.

About this task

To edit a workspace:

Procedure

1. Open the workspace that you want to edit, then click ▦ **Properties**. The Properties editor opens. The tree on the left shows the workspace name selected, with the workspace properties on the right.
2. Make any changes to the **Workspace Identity** and **Workspace Options**. You can always return to the general workspace properties by clicking the name of the workspace at the top of the Properties tree.
3. If you want to edit the properties for a particular view, select the view from the Properties tree.
4. When you are finished editing the workspace or view, do one of the following:
   - Click **Apply** to save your changes, then select another view to edit from the Properties tree.
   - Click **OK** to save your changes and close the Properties editor.
5. If you want to keep the changes you made for future work sessions, do one of the following:
   - Click 💾 **Save** to update the workspace properties with the changes you just made.
   - Click **File** > **Save Workspace As** to save this as a new workspace and leave the original workspace unchanged.

Results

Changes you make to a workspace are available only to your user ID; no other users will see your changes. The exception is when you are in workspace administration mode, where any workspace customization you do will be available to all users.

Renaming a workspace

You can rename a workspace to better suit your environment.

About this task

Use the following steps to rename a workspace:

Procedure

1. Open the workspace whose name you want to change.

2. Click ▦ **Properties**. You can also open the Workspace properties from within the Properties editor by selecting the first item in the tree, which is the Navigator item to which the workspace belongs.
3. Edit the title in the **Name** text box.

Deleting or restoring a workspace

You can delete a workspace of your creation and restore the original version of a predefined workspace or one created by the administrator.

Before you begin

👤 To delete or restore a workspace, your user ID must have Workspace Author Mode permission.

About this task

Tivoli Enterprise Portal features enable you to create and edit your own workspaces. You can also edit predefined workspaces and those created by the system administrator (unless they are protected).

Deleting a workspace
About this task

Use the following steps to delete a workspace:

Procedure
1. Open the workspace you want to delete. You cannot delete predefined workspaces. As well, you cannot delete workspaces that were created in workspace administration mode unless you are currently in workspace administration mode.
2. Click **File > Delete Workspace**. This item is disabled if the workspace is predefined, was created in administration mode and your user ID is not currently in workspace administration mode, or if the workspace properties item Do Not Allow Modifications is selected.

   Note: When you delete a workspace, all link definitions that reference it will be deleted also.
3. When a confirmation message is displayed, click **Yes** to delete the current workspace or **No** if you change your mind. The default workspace for that Navigator item displays.

Restoring the original workspace
About this task

If you have edited a predefined workspace or one created in administration mode, you can always restore the original version. To restore the original workspace, complete the following procedure:

Procedure
1. Open the predefined or administration-mode-created workspace you want to restore. If you want to retain the original and edited versions as two separate workspaces, select Save As from the File menu and give the workspace a new name, then repeat this step.

2. Click **File > Restore Original Workspace**. This item is disabled if the original predefined or administrator-created workspace is already displaying or if this is a workspace created with your user ID.

3. When a confirmation message is displayed, click **Yes** to restore the original workspace or **No** if you change your mind. The workspace reverts to its original state.

Defining links between workspaces

The Tivoli Enterprise Portal enables you to define a link from one workspace to another. Then you can quickly jump to a related workspace or a sequence of workspaces that you like to frequently review.

You can define a link from a Navigator item, or from a table, chart, graphic view within a workspace. A link from a workspace or graphic view provides a shortcut from one workspace to another, as does a link from a table row (or chart data series). Links from a table row are different in that they can use data from one or more values in the row to locate and filter the target workspace.

Your user ID requires Workspace Author Mode permission to create and maintain workspaces, including links. If the view tools are disabled, you do not have this permission.

Related tasks:

"Linking to a workspace" on page 42
Use these steps to link to a workspace that has been targeted from the current Navigator item or view.

Creating a new link

Use the link wizard to define a link from a Navigator item, graphic view, or query-based view that opens the targeted workspace in the same or new windows.

Before you begin

The workspace link feature in the Tivoli Enterprise Portal enables you to create simple links or advanced links. A simple link can originate from a Navigator item or from a data series in a chart or table. Context information is passed from the source workspace to determine which target workspace to open.

About this task

Complete these steps to create a new workspace link.

Procedure

1. Create a new link:
   a. Open the source workspace.
   b. Right-click one of the following possible link source locations:
      The highlighted Navigator item for this workspace
      An icon in a graphic view
      An object in a Tivoli Management Services Infrastructure view
      A row in a table view or situation event console view
      A data point in a chart view
   c. Click **Link To >** ✐ **Link Wizard**.
   d. Select ◉ Create a new link and click **Next>**.

e. Type a **Name** and **Description** to identify the link, then click **Next**.

2. Select the link type and target:

   a. Select one of the following link types and click **Next** to open the Workspace Link Wizard - Target Workspace window.

      ⦿ **Dynamic**

      The dynamic link method offers flexibility in your choice of workspace targets. A dynamic link is similar to a relative link and, with some link definitions, behaves in the same way. What is different is that the path to the target workspace uses a system identifier passed by the source query-based view. The identifier can be managed system name, IP address, host name, Systems Management Facilities ID, or Product Code.

      ○ **Absolute**

      An absolute link stores the managed system name of the chosen target, and goes only there. For example, if you target the **Databases Information** workspace on the **MS SQL Server** branch for the **Paris** computer, that is the workspace that will open. The **Database Information** workspace for the **Geneva** computer is not considered.

      ○ **Relative**

      A relative link stores the *type* of source workspace and target workspace. The resulting link is available from any workspace that is of the same type. If more than one workspace fits the definition of the target, you are presented with a list of workspaces to choose from. For example, a link built from the Paris computer's **Windows OS** workspace to its **Memory** workspace will also be available from the Geneva computer's **Windows OS** workspace and link to its **Memory** workspace.If you are linking between Navigator views, relative linking is not recommended. Instead, use the absolute or dynamic method for links to the Navigator Logical view and other custom Navigator views.

   b. If you are targeting a workspace in a different Navigator view, select the view from the **Navigator View** ▾ list.

   c. In the **Navigator** tree, expand the branches and select the Navigator item for the workspace you want to target. The Navigator view is presented in the left frame. As you open branches and select Navigator items, the right frame show their workspaces. With an absolute link, the path you select to the target is interpreted as an exact match, whereas, with a relative or dynamic link the path is interpreted as the pattern to follow from the link source to the target workspace.

   d. In the **Workspace** pane, select the workspace to target; or, if this is a dynamic link and you want to more generally classify the link, select ☑ **This link does not target a specific workspace**.

   e. If you would like the workspace to open in a new window instead of replacing the current workspace, select ☑ **Always open target workspace in new window**. If you are running the browser client and your browser supports tabbed pages, the workspace will open in a new tab.

   f. Click **Next** or, if this is an absolute or relative link and you are finished defining the link target, click **Next** twice to see the the summary.

3. If this is an absolute or relative link and you want to refine the link definition, click **Next** and add the expressions.

4. If this is a dynamic link, the Target Filters window is displayed for you to select the type of identifier for the target and specify the value to use:

   a. Click the identifier to use: Managed system name, Hostname (distributed systems only), IP address (distributed systems only), or SMFID (z/OS-based systems only) If you are not sure which one to choose, click **Managed System name**.

   b. Click **Modify Expression** to open the Expression Editor, then click **Symbol**. The list is a reverse hierarchy of available ▼ symbols starting from the source context and the current Navigator item and ending at the root Navigator item.

   c. Select a ▼ symbol that provides the value for the system identifier. For example, **Server Name** will provide the value needed for **Managed system name** when the link is launched. You might need to open different branches to locate a symbol for the type of target filter chosen.

      ▦ Selected Row
         📁 Attributes
            ▼ % Usage
              ▼ Current Size
                ▼ Server Name

   d. Click **OK** add the detailed name for the attribute to the Expression Editor. $kfw.TableRow:ATTRIBUTE.NTLOGINFO.ORIGINNODE$

   e. Click **Evaluate** to verify that the variable resolves correctly. The Value shown must be valid for the identifier type you specified. For example, for a computer named TIVOLIUSER, Primary:TIVOLIUSER:NT is a valid Windows OS managed system name and TIVOLIUSER is a valid host name.

   f. If no value is displayed or you get a syntax error or a value that does not identify the system, click **Clear** and select a different symbol or try another identifier.

   g. In addition to **Symbol**, you can click **Operator** and **Function** to find and select other elements to complete the expression. When you are finished with the expression, click **OK** to insert the expression in the current target filter row. Note that expressions can evaluate to patterns using the asterisk (*) and question mark (?) wildcard characters. This allows for pattern matching filters.

   h. If the target workspace was updated or is new for the latest version of your monitoring product and your environment has a mix of versions, you can disable ☑ **Allow this link to target any version of the target workspace** or leave it enabled:

      ☑ Enabled is the default setting. The target workspace can be chosen for any version of the monitoring agent. If the query for a view specifies attributes that are available only in the newer version and the target workspace is for an older version of the monitoring agent, no values will be returned for those attributes. ☐ Disable the option if you want the link not to be offered for target workspaces on earlier versions of the monitoring agent.

   i. Click **Next** and refine the link definition; or click **Next** twice and review the summary.

5. In the Workspace Link Wizard - Summary page, review the summary, then click **Finish** to save the link definition, **Back** to return to a previous window to make additional changes, or **Cancel** to close the Link Wizard without saving your changes.

What to do next

If the target workspace is dependent on information that is passed from the source workspace (users will get an error message if they attempt to open the workspace from the list of available workspace), then edit the workspace properties to hide the workspace unless it is targeted by a link definition.

Testing the link

You can test the link you have created to ensure that it works correctly.

About this task

To test a link you have created:

Procedure

1. In the source workspace, right-click the location where you created the link.
2. Click **Link To** and select the link from the ✐ list. The value of the symbols specified in the link definition are used to select the target workspace and filter its content. An example is a link from a row in a view that includes the **Process Name**, and the value of Process Name is passed to the target workspace. The view whose symbol was targeted will show information about that particular process.
3. If the target workspace does not display what you expected or a Target not found message is displayed, edit the link.

   Tip: If you were prompted to select the target workspace but wanted the link definition to be specific so you would not have to choose from a list, do one of two things: If you had chosen Relative as the link type, define a new link and select a link type of Absolute; or make it Dynamic and supply the expression in the target filter to target the specific workspace.

Editing a link definition

You can edit a link definition after it has been created.

Procedure

1. Open the workspace where the link originates (source workspace).
2. Right-click the link origination point, which can be one of the following:
   * The highlighted Navigator item for this workspace
   * An icon in the graphic view
   * A row in the table view
   * A row in the situation event console view
   * A bar in the bar chart view
   * A slice in the pie chart view
   * A point in the plot chart view
3. Click **Link To >** ✐ **Link Wizard**.
4. In the Workspace Link Wizard - Welcome page, select ◉ **Modify an existing link** and click **Next>**.
5. In the Workspace Link Wizard - Link to Modify page, select the name of the link to edit, then click **Next>**. Before clicking **Next**, you can clear the ☐ **Hide links that are not available for the select view** check box to see all the links for this workspace.

6. Edit the name or description, then click **Next>** to open the Link editor.
7. Edit the link.
8. Click **Next** until you reach the summary page, then click **Finish**.

Deleting a link definition

You can delete a link definition when it is no longer useful.

About this task

Follow these steps to permanently remove a link definition.

Procedure

1. Open the source workspace.
2. Right-click the link origination point, which can be one of the following:
   - The highlighted Navigator item for this workspace
   - An icon in a graphic view
   - An object in a TMS Infrastructure view
   - A row in a table view or situation event console view
   - A bar in a bar chart view
   - A slice in a pie chart view
   - A point in a plot chart view
3. Click **Link To >** ✎ **Link Wizard**.
4. Select ◉ **Delete one or more existing links** and click **Next**.
5. Select the name of the link to delete. You can select additional links with Ctrl+Click or select all links from the first selection to this one with Shift+Click.
6. Click **Next** until you reach the summary page, then click **Finish**

   Note: When you delete a workspace, all link definitions that originate from it will be deleted also. This is also true for any absolute links that target this workspace.

Adding or editing a link anchor

You can add or edit a link anchor by using the Link Anchor Properties editor to hide the link indicator, disable it for table rows that do not meet the link criteria, and to assign a different default link or no default link.

About this task

When you have a link that originates from a table row or a graphic view icon, a link indicator is displayed: Click ✎ to open the default link or to see and select from a list of possible links.

Procedure

1. Open the workspace where the link originates (source workspace).
2. Right-click a table row or the graphic view icon where you originated the link.
3. Click ✎ **Link Anchor**.
4. In the Link Anchor Properties editor, select a default link from the list or select **(No Default)**. When you select the option to show the link indicator (next step), you can click the indicator to go directly to the default linked workspace or, if no default link was assigned, to see and select from the list of available links.

5. Clear ☑ **Show Link Indicator** if you want no 🖉 link indicator in the first column of the table or overlaying the graphic view icon. If you disable the indicator, users can still link to the target workspace by right-clicking a table row or the linked graphic view icon.

6. Clear ☑ **Link Indicator Always Enabled** only if this is an advanced link from a table, with filtering, and you want to enable the indicator only for rows where the link is available. An example would be a link that is not available for the *Idle* process. The link indicator in this row would be 🖉 dimmed to show that it was disabled.

7. Click **OK** to close the window, then click 🖫 **Save** to save your changes to the workspace.

Link examples

Review these link examples for ideas on how to use the Link Wizard to build links between workspaces.

The graphic shows where you can start the Link Wizard to define a link, and from where you can then launch the link to the workspace that was targeted by the link definition. The mouse icons show how the link is launched: click a link anchor; or right-click a Navigator item or chart view data item. A link that originates from a Navigator item or graphic view is often a simple link to the target workspace, whereas a link from a table row or chart point might be to a workspace with more details about that row. Notice that the link anchors showing in the table view and graphic view appear automatically when you define a link from these views.
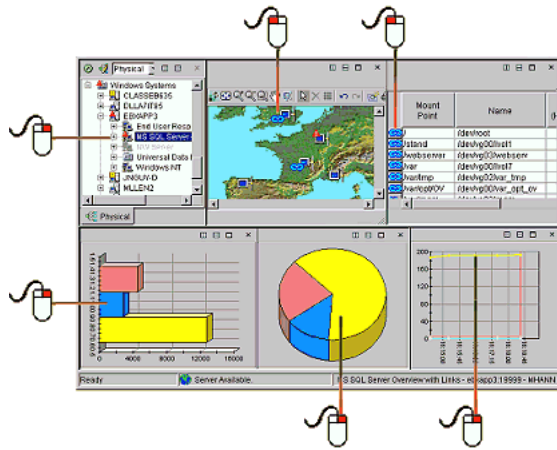
Dynamic linking

The dynamic link type has you specify an identifier for the target workspace. When the source is a table row or chart point, you can often select a symbol for an attribute that identifies the managed system.

Target Filter:
    Managed system name

Expression:
    $EventConsoleRow:ATTRIBUTE.ISITSTSH.ORIGINNODE$

Symbol:
    📁 Values
        📁 Link
        📁 Selected Row

📁 Attributes
  ᵥ Age
  ᵥ Display Item
  ᵥ Global Timestamp
  ᵥ Impact
  ᵥ Node
  ᵥ Origin Node
  ᵥ Severity
  ᵥ Status
ᵥ Id
ᵥ Name
ᵥ Type
🔲 Situation Event Console
🔲 Query
🔲 Workspace
📁 Node
📁 System
🔲 Root
🔲 Global Symbols

When you cannot find a symbol for the managed system, consider using another identifier. In the example here, the link source is a Navigator item and the target filter chosen is either the Hostname or IP address.

Target Filter:	**Target Filter:**
Hostname	IP address
Expression:	**Expression:**
$NODE:-1022$	$NODE:-1022$
Symbol:	**Symbol:**
📁 Values	📁 Values
📁 Link	📁 Link
🔲 Tree - Physical	🔲 Tree - Physical
ᵥ Id	ᵥ Id
ᵥ Name	ᵥ Name
ᵥ Type	ᵥ Type
🔲 Query	🔲 Query
🔲 Workspace	🔲 Workspace
🔲 Tree - Physical	🔲 Tree - Physical
🔲 Query	🔲 Query
🔲 Workspace	🔲 Workspace
📁 Node	📁 Node
ᵥ Id	ᵥ Id
ᵥ Name	ᵥ Name
ᵥ Network Address	ᵥ Network Address
ᵥ Type	ᵥ Type
📁 System	📁 System
🔲 Root	🔲 Root
🔲 Global Systems	🔲 Global Systems

Relative link from a Navigator item

You can link from a workspace at the 📁 system level of the Navigator to a workspace at the 📁 platform level. Because it is relative, the link is available from any system level Navigator item (NODE1, NODE2, NODE3) to the parent Windows Systems workspace. In an environment with many Windows systems being monitored, this link enables a quick jump from the expanded branch to the parent workspace.

        🏢 Enterprise
                📁 Windows Systems
                        📁 NODE1
                        📁 NODE2
                        📁 NODE3
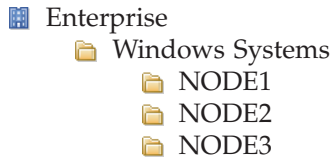
Source point:
> Any of the NODE Navigator items

Link Type:
> Relative

Target Workspace:
> 📁 Windows Systems > Windows Systems

Parameters:
> The expression is built automatically from the source context and the workspace that was selected as the target.
>
> - 🔣 contextIsAvailable:
> - $kfw.TreeAdapter:-1021A$ == $kfw.LinkSymbolsGroup:PBASED.SOURCE_CONTEXT.0.-1021A$ && $kfw.Context:-5000$ &&
> - $kfw.TreeAdapter:-1021A$ == $kfw.LinkSymbolsGroup:PBASED.SOURCE_CONTEXT.3.-1021A$ && $kfw.Context:-5000$
> - 🔣 linkIsEnabled: true

linkIsEnabled usage example

Follow along with these steps to see how a link to the System workspace was created using the 🔣 **linkIsEnabled** symbol. This example tests the value of the Product cell in the selected row to determine whether or not to offer the link on that row.

1. Open the Managed System Status workspace in the Navigator Physical view: Click 🏢 **Enterprise** to open the default workspace; right-click 🏢 **Enterprise** and click **Workspace > Managed System Status**.
2. Right-click a row in the Managed System Status table and click **Link To >**✏️ **Link Wizard**.
3. Create a new link, name it `System Details`, select a link type of **Relative**, and target the 🗔 Windows Systems workspace.
4. Click **Next** to go to the Workspace Link Wizard - Parameters.
5. Click 🔣 **thelinkIsEnabled** row. The Expression box reads `true`.
6. Click **Modify Expression**.
7. In the Expression Editor, click **Clear**, and then click **Symbol**.
8. In the **Symbols** list, select 🔣 **Product** and click **OK**. The expression is displayed as `$kfw.INODESTS:ATTRIBUTE.INODESTS.PRODUCT$`.

9. Click **Operator**, select ▣ **Equal** and click **OK** to see the == operator.

10. Type "NT" at the end of the expression, which should then read,
    `$kfw.INODESTS:ATTRIBUTE.INODESTS.PRODUCT$ == "NT"`.

11. Click **Evaluate**. If the row from where you started the Link Wizard has `NT` in the Product column, the test will be `true`; otherwise, it will be `false`.

12. Click **Finish**.

Click the ✎ link indicator in one of the NT rows to open the Services workspace for that monitoring agent. Note that you can disable the link indicator for invalid rows by clearing ☐ **Link Indicator Always Enabled**.

Text in link expressions

You can add text and spaces to the link expression by enclosing them in " double quotes and joining to any other functions with a + plus sign. Here are some examples.

Link expression	Resolves to
`$kfw.LinkSymbolsGroup:PBASED.SOURCE_CONTEXT.2.-1022$`	`Windows Systems`
`"My" +$kfw.LinkSymbolsGroup:PBASED.SOURCE_CONTEXT.2.-1022$`	`My Windows Systems`
`"My "+$kfw.LinkSymbolsGroup:PBASED.SOURCE_CONTEXT.2.-1022$ + "Now"`	My Windows Systems Now
`$kfw.LinkSymbolsGroup:PBASED.SOURCE_CONTEXT.2.-1022$ +" in the Navigator "+$kfw.TreeAdapter:-1022$ +" view"`	`Windows Systems in the Navigator Physical view`

Symbols in text

You can reference a variable (symbol) in a table or chart view header or footer, or in a notepad view, and have it resolve to a value when you link to the workspace. The symbol value is defined in the link wizard customization page.

Link anchor properties

The link anchor properties enable you to show a visual indicator on tables where customized links have been defined, and to establish a default link that opens when the user clicks the indicator.

Default Link
> When you select the option to show the link indicator, the user can click the ✎ link indicator to go directly to the default linked workspace or, if no default link was assigned, to see and select from the list of available linked workspaces. Select from the list of target workspaces or select **(No Default)** to have no default link.

Options
> ☑ **Show Link Indicator** adds a new first column to the table, within each row; or it overlays the Navigator item icon in the graphic view with ☑ . Clear this option if you want no link indicator in the table or graphic view.

> ☑ **Link Indicator Always Enabled** shows the links in every row of the table or over every linked Navigator item icon in the graphic view, which is normally what you want. Some links, however, are conditional, based on the value of a particular cell (or a derived value). For example, a link might be defined to work only on rows that have a thread count value over 30. You would clear this option to have the link indicator display only on the rows where the thread count is over 30. The advantage of this

option is it guides the user to the valid link rows; the disadvantage is that it takes more time for the table to render.

Chapter 6. Customizing workspace views

Tivoli Enterprise Portal workspaces can be customized to show different views.

The Tivoli Enterprise Portal workspaces can be composed of more than a dozen different views. There are query-based views for displaying the attribute values gathered by monitoring agents, event views for displaying events and their status, and a number of special purpose views.

Creating a new view

Create a new view to customize a workspace to show the reports and other information that give you a comprehensive view of your environment or particular aspects of it.

You can add as many views to a workspace as you can easily see within the confines of the window. Creating a new view is a matter of dividing one view into two and editing one of them.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new space, click ▥ **Split Vertically** or ▤ **Split Horizontally** in an existing view. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty. If you want to replace a view with a different type of view, skip this step.
3. Click the tool for the type of view you want: ▦ ◉ ▥ ▩ ▩ ◕▧ 📄 📋 📝 📠📟 ◪ ⛁ 🎴 🎴 🎴 . The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
4. Click inside the space to adopt the new view. If this is a table or chart view and no query has been chosen, or the query is inappropriate for the type of chart, no data is displayed. You might be prompted to assign a query.
5. Some views require more information:
   - If the Select Attribute window opens, you replaced a table view with a chart view, which can show only numeric attributes. Select an attribute (use Ctrl+Click or Shift+Click to select more), then click **OK**.
   - If you are adding a terminal view, enter the host name and port number in the Terminal Emulator Configuration window.
6. Select ✎ **Edit properties** to open the Properties editor to change the content or format. The tabs displayed depend on the type of view. The Tivoli Enterprise Console event viewer and the take action view have no associated properties.
7. As you edit the properties, click **Test** to see your changes in the **Preview** area at the top of the window, **Apply** to save your changes and keep the window open for selecting another view to edit, or **OK** to save your changes and close the window.

Attention: The **Preview** area is showing a sample of what your view will look like in the associated workspace. This area is restricted by space constraints and some labels, identifiers, and other view-level adornments might be hidden, truncated, or partially rendered. Your view will be rendered correctly when applied to the associated workspace where space constraints are not a concern.

8. If you want to keep the view in this workspace for future work sessions, do one of the following:

   - Click 💾 **Save** to update the workspace properties with the new view.
   - Select **Save Workspace As** from the **File** menu to save this as a new workspace and leave the original workspace as it was.

   Note: If no information appears in a table or chart, it might be because no data samplings meet the query or filter criteria. Regardless of the managed systems assigned to a query, a gauge chart (circular or linear) can show data from only one managed system. If you are using the query in a workspace where multiple systems are applicable (at the platform level of the Navigator physical view), the data shown will be from the first managed system in the Assigned list of the Query Results Source tab.

Editing a view

You can change the style or other aspects of a view through the Properties editor. The Properties editor is associated with the current workspace.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. In the view title bar, click 🖉 **Edit properties**. The Properties editor for the workspace opens and displays the properties of the selected view.
2. Edit the view properties using the following tabs:

Option	Description
🗊 Query	Specifies what data must go in the chart or table or table-based topology view. See "Query" on page 92.
🔁 Filters	Refines the query-based view by filtering out unwanted data from the chart or table. See "Filters" on page 93.
🔳 Thresholds	Establishes threshold values and color indicators for the table view. See "Thresholds" on page 95.
📄 Configuration	Specifies the script to run or the connection to make whenever you open the terminal view. See "Configuration" on page 96.
A𝒶 Style	Changes the appearance of the view. See "Style" on page 96.

View type	Available Properties tabs
▦ Table	Query, Filters, Thresholds, Style

View type	Available Properties tabs
Charts: ⬤ Pie, ▉ Bar, ▦ Plot, ▨ Area, ⬤ Circular gauge, ▮ Linear gauge	Query, Filters, Style
📋 Notepad	Style
📄 Message log	Filters, Thresholds, Style
🗔 Universal message console	Filters, Thresholds, Style
🖼 Graphic	Style
🖥 Terminal	Query, Configuration, Style
🗔 Browser	Style
🔗 Topology	Style
▦ Situation event console	Filters, Thresholds, Style
🗔 Common Event Console	Filters, Thresholds, Style

While editing the properties, click **Test** to see your changes in the **Preview** area at the top of the window.

Attention: The **Preview** area is showing a sample of what your view will look like in the associated workspace. This area is restricted by space constraints and some labels, identifiers, and other view-level adornments might be hidden, truncated, or partially rendered. Your view will be rendered correctly when applied to the associated workspace where space constraints are not a concern.

In the Filters and Thresholds tabs, if you need to add or remove rows or clear cell contents, right-click and select one of the pop-up menu options.

3. When you are finished editing the view, click **Apply** to save your changes, then select another view to edit from the Properties tree; -OR- click **OK** to save your changes and close the window.

4. To keep your changes, do one of the following:

   - Click 💾 **Save** to update the workspace properties with the changes you just made.

   - Select **Save Workspace As** from the **File** menu to save this as a new workspace and leave the original workspace unchanged.

What to do next

You can also resize a view by dragging the borders, and swap locations with another view by dragging it by its title bar to the new position. You cannot move views that are part of a workspace that is protected from editing, which applies to predefined workspaces.

Query-based views

Query-based views are the table view, chart views, and *relational table-based* topology view. They are distinguished from other workspace views by their ability to show monitoring data from agents and from external JDBC or ODBC data sources for which you have written SQL statements. Use these views for on-demand reporting of current and past (historical) data samplings.

The predefined workspaces for your product include chart and table views. These views are the first step to getting something meaningful from the data being collected. When you understand what values and states are causing problems, you

can refine your views to show what is important. Additionally, the topology view can express the relationships between attributes as connected objects when the topology source is relational table-based.

Table A table shows a column for each attribute and one or more rows of data. You can also get table views to show data samplings over a period of time. The attribute group queried by the view must be either a historical attribute group or have been configured for historical data collection, in which case you will see a Time Span tool in the view toolbar for setting the time period.

Chart A pie chart has a slice for every data point in a single data series (row). Pie charts best show the proportional value of related attributes to a whole, such as the percentage attributes that show how memory is being used.

  A bar chart displays a bar for each data point. Bar charts are best suited for comparing values among related attributes. The stacking bar chart best shows multiple values for the same attribute.

  A plot chart shows changes over a period of time by drawing a continuous line from one data point to the next, one data point for each data sampling and one line for each attribute selected. The plot chart best shows trends over time and among related attributes.

  An area chart is almost identical to a plot chart except that the area between the axis and each plot line is filled with color or a pattern.

  A circular gauge chart shows the proportional amount of a data series, one gauge for each attribute chosen. This type of chart is well suited for showing individual elements that change frequently, such as percentage user time.

  A linear gauge chart shows the collective value of every item in a single data series, one gauge for each attribute chosen. This type of chart is well suited for showing cumulative values.

Topology
  A Topology view shows the arrangement of monitored components associated with its Navigator item. It is considered a query-based view when it has been created as a relational table-based topology view, which is appropriate for attributes that can show the relationship to other attributes in their group as connected objects.

Selecting a query for a view

Query-based views – chart, table or relational table-based topology – must have a query assigned to specify the data to retrieve.

Before you begin

If you do not see the Queries tool, your user ID does not have View or Modify Query permissions. If you can see the tool but it is disabled, your user ID does not have Workspace Author Mode permission.

About this task

If a newly created chart, table or relational table-based topology view has no query assigned to it, you must select one. If the view shows data, but you want to show different attributes, change the query.

Procedure

1. In the view title bar, click ✎ **Edit properties**. The Properties editor for the workspace opens and displays the properties of the selected view.

2. In the **Query** tab, click ⬚ **Click here to assign a query** to open the Query editor. The Query editor shows a tree on the left organized by Tivoli monitoring product and then by attribute group.

3. Click ⊞ to expand the ⊙ Monitoring agents and ⬚ attribute group folders to see the queries.

4. Select a ⬚ Query title to see its description. The query description, edit date, and specification are displayed in the right frame. The selected options indicate which attributes from the group are included in the query. The cells in the numbered rows show any filter criteria for the query.

5. Optional for Tivoli monitoring product queries: To apply the query to specific managed systems rather than all managed systems of this type, do the following:

   a. Click the **Query Results Source** tab.

   b. Select ◉ **Let user assign explicitly**.

   c. Remove from the Assigned list those systems you do not want in the view.

6. Click **OK** to select the query and return to the Properties editor.

Results

The Description and Formula areas of the Query tab show the name, description, and formula of the query that you selected; the **Preview** area shows how the rendered view will look when the query is applied.

Attention: The **Preview** area is showing a sample of what your view will look like in the associated workspace. This area is restricted by space constraints and some labels, identifiers, and other view-level adornments might be hidden, truncated, or partially rendered. Your view will be rendered correctly when applied to the associated workspace where space constraints are not a concern.

What to do next

Make any other adjustments to the view properties. After you click **Apply** or **OK** in the Properties editor, the query is applied to the view and the results shown in the workspace.

If you chose a custom SQL query and it includes a variable in the statement, you are prompted to enter a value after you click **Apply** or **OK**. The value filters the data that is displayed in the view, but is temporary and cannot be saved with the workspace.

Regardless of the managed systems assigned to a Navigator item, a gauge chart can show data from only one managed system. If you are using the query in a workspace where multiple managed systems are applicable (at the platform level of the Navigator Physical view), the data shown is from the first managed system in the Assigned list.

Note: If you change the query for an existing chart or gauge view, you must ensure that you select which columns you want to chart using the **Filter** tab in the view's **Properties** dialog. By default, no columns are selected for charting.

Filtering a query-based view

You can filter the information displayed on the view to show specific attributes and rows of data.

About this task

Use the ⚡ **Filters** tab to control the display to show only those attributes and rows of data that are of interest to you. Using this method, you can easily see and compare the relevant information in a small viewing area.

If you are filtering a view that returns so many rows that response time is compromised, consider using a query to apply a pre-filter that reduces the amount of data to retrieve.

Procedure

1. If the Properties editor is not open, right-click the table view and select Properties from the pop-up menu.
2. Click the ⚡ **Filters** tab.
3. Select a ☑ check box beneath a column heading to include the attribute in the view; clear a ☐ check box to exclude it. If you clear a check box, you can still add a filter expression for that attribute (next step).
4. If the attributes report multiple rows of data and you want to limit what displays, enter a filter expression:
   - Click a cell under a column to activate the editor and formula bar. Then click inside one of the fields to:

     ⟦*v*⟧ change the function from the default **Value of expression** (or **Compare date/time** for time attributes) to another function.

     ⟦==⟧ change the relational operator from the default Equal to Not equal, Greater than, Greater than or equal, Less than, or Less than or equal.

     ⟦▭⟧ enter the value to be compared with the values arriving from the managed system. Enumerated attribute values, such as *online* or *offline*, will have a list of possible choices.
5. Repeat steps 3 and 4 if the filter criteria involves multiple columns or multiple rows:
   - Expressions entered in the same row have an AND relationship. In the example below, a row will be written when the Process Name includes `java` -AND- user time exceeds 5%. (Scan for string within a string equals java for Process Name AND % User Time < 5.)

	Process Name	% User Time
1	☑	☑
2	🔤 == 'java'	>5
3		

   - Expressions in different rows have an OR relationship. In the example below, a row will be written when the value of Process Name is either *System* or *Services*.

	Process Name	% User Time
1	☑	☑
2	== 'System'	

	Process Name	% User Time
3	== 'Services'	

6. Do one of the following:
   - Add icons or background colors to a table view when thresholds have been reached.
   - Change the title of the view or apply other styling.
   - Apply advanced filtering to show only the first (highest value) or last (lowest value) data rows, such as the first 25 rows.
   - To edit the properties of another view in this workspace, click **Apply** to save your changes, then select the next view from the Properties tree.
   - To apply your changes and close the Properties editor, click **OK**.

Properties editor

Use the properties editor to change the general properties of the workspace and to customize the appearance and behavior of the views in that workspace. Every workspace has general properties and an additional set of properties for each view it contains.

Clicking ⬚ **Properties** in the main toolbar or ✎ **Edit properties** in the view's title bar opens the Properties editor. The Properties tree shows the name of the workspace at the top of the hierarchy, followed by a Views folder, then folders for every type of view in the workspace. Appearing to the right of the tree are the workspace properties or view properties, depending on which is selected in the tree.

Workspace properties

Every workspace has a name and general characteristics defined in the workspace properties.

When you save a new workspace with **File → Save Workspace As**, click ⬚ **Properties**, or, if the Properties editor is already open, click the top item in the tree to see the workspace properties.

Workspace Identity

Name   The title of the workspace. The workspace name is displayed in the title bar or browser mode banner, in the **View → Workspace** list, or in the **Link To** list menu, depending on how you have set the workspace options (see below). Do not use an & ampersand or # pound sign in the title.

Description
Optional. The description can be a full text description of the workspace.

Workspace Options

☐ **Assign as default for this Navigator Item**
Makes this workspace the default for the Navigator item (the one that opens when you click the Navigator item).

☐ **Assign as Home Workspace**
Specifies this as the workspace that opens whenever you start Tivoli Enterprise Portal, much like your browser home page.

☐ **Do not allow modifications**

Protects the workspace from permanent changes and must be selected when you are editing in Workspace administration mode. Users who have access to this workspace can change it temporarily, but if they attempt to ▣ save, the Save Workspace As window opens so that it is saved as a new workspace with a different name. Also, when this item is enabled you cannot delete the workspace.

☐ **Only selectable as the target of a Workspace Link**

Marks the workspace as hidden except when it is a possible Link To destination. The workspace does not display in the list of available workspaces for a Navigator item.

☐ **Always open workspace in new window**

When enabled, this option causes the workspace to open in a new window, or in a new tab if your browser supports tabbed pages.

View properties

Edit the properties of a workspace view to control its appearance and behavior.

The view properties are displayed when you click ✎ **Edit properties** in the view's title bar or when you click the view in the **Properties** tree in the left pane of the Properties editor. The view properties open in the right frame with one or more of these tabs, depending on the view type, such as the ▣ **Thresholds** tab for the table view.

Preview is a sample of the view before you selected it in the Properties editor. It gives you a preview of how the view will look when you save (with **Apply** or **OK**) your changes. Preview is typically used with query-based views whenever you select a different query or edit the query, define a filter or threshold, or modify the style.

Attention: The **Preview** area is showing a sample of what your view will look like in the associated workspace. This area is restricted by space constraints and some labels, identifiers, and other view-level adornments might be hidden, truncated, or partially rendered. Your view will be rendered correctly when applied to the associated workspace where space constraints are not a concern.

The editor has buttons to **Test** your changes in the Preview area to see how they will look, **Apply** your changes to the view to save them now, and exit with (**OK**) or without (**Cancel**) saving your changes. If you plan to edit more than one view while in the Properties editor, click **Apply** to save your changes to the first view before editing the next.

Query:

A query specifies the attribute values to display in the table view, chart views, relational table-based topology view, and, optionally, the terminal view.

As well as queries of attributes from your managed systems, you can write SQL queries for an ODBC-compliant database. After the data is retrieved, you can filter out any unwanted data from the view.

Multiple views in the same workspace can use the same data retrieved by the query, use their style to format the data, and filters to control what is shown. Using the same query for multiple views in a workspace conserves resources at the

portal server and the monitoring server. Also, using a query that includes only the attributes you want to see in a view saves resources.

These views have a **Query** tab where you can select what attribute group and, when there are multiple queries for the same group, which attributes to populate the view with.

📖 **Click here to assign a query.**

> Opens the Query editor so you can select an existing query or create a new one. The queries available depend on where the workspace is in relation to the Navigator hierarchy. The exception is Tivoli Enterprise Monitoring Server: these are queries of the common attributes, which are available for all workspaces.

Description

> **Name** of the query.
>
> **Description** of the query.

Formula

> Shows the formula for the query. Note that any column function added to the query is not displayed here.

View-level Page Size

> Controls the page-break size of the table, message log, pie chart or bar chart view. For example, a pie chart gives you one pie for each row of data returned, so 50 rows will give you 50 pies in the view. By limiting the page size to three rows, you can have three pies on a page, using ⬇ to turn to the next page. The smaller page size speeds up data retrieval to the workspace because only one page is retrieved at a time rather than the entire row set.

◉ **Use default** retrieves 100 rows of data per page.

○ **Return all rows** retrieves the entire row set to the view at one time. If there are many rows, the view might take some time to display.

○ **Number of rows to return** ☐ enables you to specify the exact number of rows to apply to a page. Keep the number to a size than can be retrieved in a reasonable period of time.

Filters:

Chart, table, message log and event console views have a Filters tab so you can select which attributes and what data to show.

Note that you can also use a query to filter the data. The difference is that the query filters the data before it is retrieved and the view filters the data after it has been retrieved.

Filters

> All attributes associated with the query are displayed as column headings. Use the check boxes to add or remove attributes in the view, and use the editor to add filters.

☑ Select the check box beneath a column heading to include the attribute in the view; or clear the box to exclude it. Even if an attribute is excluded from the view, you can still write a filter expression for it.

Chart views can filter on numeric attributes only; the check box for text and timestamp attributes is disabled ▢ .

Click a cell under a column to activate the editor and formula bar. The active cell has three fields 〔 ν 〕〔 == 〕〔　　〕 :

- 〔 ν 〕 Function shows the function chosen for the attribute, which is set to *Value of expression* by default. Click the button to open and select from a list of possible functions, which vary depending on the type of attribute. For example, the available functions for text are *Scan for string within a string* , *Return a subset of the string* , and *Value of expression* . Numeric attributes have *See if a value is within a range* , which checks if the value falls within the range you specify.

- 〔 == 〕 Operator shows the relational operator, which is == Equal by default. Click the button to open and select from a list of operators: == Equal, != Not equal. > Greater than, >= Greater than or equal, < Less than, and <= Less than or equal.

- 〔　　〕 Text box shows the value of the attribute to be compared with the values arriving from the agent. Click inside the cell to edit the value. For enumerated attribute values, such as "online" or "offline," you see a list box to click for the possible choices.

You can specify many expressions to form a filter:

Expressions in the same row have an AND relationship; expressions in different rows have an OR relationship. The example shows both: A row will be written to the chart or table if the process name includes "java" AND the processor time is more than 30% OR the user time is more than 20%.

	Process Name	% Processor Time	% User Time
1	☑	☑	☑
2	abc == 'java'	> 30	
3	abc == 'java'		> 20

Show Formula opens a graphical view of the filter.

Data Snapshot

Shows the attributes returned from the query and a sampling of their values.

Set advanced filters
Opens the Advanced Filters window for setting a sort order for the retrieved values and the highest or lowest values to return.

Tip: When filtering is established through the **Filters** tab, the filtering occurs after the data is retrieved from the managed system. If many rows are returned, it might take some time to retrieve all of them and complete the filtering task. If you create a query that filters the data, then assign that query to the view, the filtering is done at the portal server, which improves performance time.

Thresholds:

The properties for table views include the Thresholds tab for defining the values or value ranges for coloring table cells ![57] or applying a severity icon ![54] .

In addition to the severity for threshold cells, you can sort the table by the severity level of a column, and you can move the mouse pointer over a cell to see the threshold criteria.

Thresholds

The threshold editor shows the attributes for the query as column headings. The cells in the first two columns are for selecting a severity indicator: The Use Icons ☑ check box, when selected, will apply the icon that is selected in the next (Threshold) column to the threshold cell when the expression in this row is true. The Threshold column lists the possible severities.

The threshold expression for the chosen severity is written in the same row. Clicking a cell activathanks. tes the formula editor and the formula bar. The active cell has three fields for completing the threshold expression ![v ==] :

- ![v] Function shows the function chosen for the attribute, which is set to *Value of expression* by default. Click the button to open and select from a list of possible functions, which vary depending on the type of attribute. For example, the available functions for text are *Scan for string within a string* , *Return a subset of the string* , and *Value of expression* . Numeric attributes have *See if a value is within a range* , which checks if the value falls within the range you specify.

- ![==] Operator shows the relational operator, which is == Equal by default. Click the button to open and select from a list of operators: == Equal, != Not equal. > Greater than, >= Greater than or equal, < Less than, and <= Less than or equal.

- ![ ] Text box shows the value of the attribute to be compared with the values arriving from the managed system. Click inside the cell to edit the value. For enumerated attribute values, such as "online" or "offline," you see a list box to click for the possible choices.

Note: For complex expressions, such as when you need to repeat an attribute name, you can double-click in a value text box to see and select an available attribute from the Select Value window. Clear the box before double-clicking or the value will be inserted between existing text and the formula syntax will be incorrect.

You can specify multiple threshold expressions to form a filter:
- The threshold expressions you enter in the same row have an AND relationship. When a retrieved value satisfies all the expressions in a row, the cell background will be colored.

- The threshold expressions in each row are independent. If a retrieved value satisfies the conditions in more than one row, the cell background will show the color for the highest severity, the lowest being Unknown and highest being Fatal. Example:

	Use Icons		Process Name	ID Process	% User Time
1	☐	🔴 Critical			> 90
2	☑	⚠ Minor			> 50

	Use Icons		Process Name	ID Process	% User Time
3	☑	✅ Harmless	abc == 'java'		> 5

⚙ Show Formula opens a graphical view of the thresholds.

Data Snapshot

> This shows the attributes returned from the query and a sampling of their values.

Configuration:

Use the Configuration tab of the Properties editor to edit the terminal view script or to specify a script to run whenever you start a terminal session in a workspace.

The properties for the terminal view have a **Configuration** tab so you can edit scripts, specify a script to run whenever you start a terminal session, and specify the host to connect to. Whenever you open the workspace, the script you specify will run automatically, and the host you specify will be connected automatically.

Configuration

> These options are the same as those given in the Emulator Scripts window that opens when you click the Manage Scripts tool in the terminal view.

Connection

> This is the same as the Terminal Emulator Configuration window that opens when you click the Connection tool in the terminal view, or add a terminal view.

Disconnect options

> ◉ **Do not disconnect when leaving workspace** keeps the session open as long as you are logged on to Tivoli Enterprise Portal, so you can return to this workspace and the session will be where you left it last.

> ◯ **Disconnect after this many minutes:** ⬜ will disconnect the works session the number of minutes entered here after you close the workspace.

Style:

Every view type except Tivoli Enterprise Console and Take Action has style properties that you can edit to affect the appearance and behavior of the view.

The 🅰 **Style** tab of the Properties editor has a thumbnail graphic with the style elements as clickable hot spots. When you click a hot spot, such as **Plot Area** for a chart view, the style settings change.

Options

> These vary for the type of view.

Title

> **Text** is the title that is displayed at the top of the table, chart, or specified element therein (such as the header). If this is a chart or table, you can also add an attribute symbol to the title that will be replaced by the value of that attribute when the view is displayed.

Note: When the workspace is opened, the header and footer are evaluated for any expressions. If you happen to have a relational operator (==, !=, >, >=, <, <=) in the text, the operator will be replaced by TRUE or FALSE unless you enclose it in quotes.

Font name is the font in which the title will be displayed.

Size is the font size in points.

Style is the format of the font. The choices vary depending on the font, but are usually Bold, Plain and Italic.

Text

Font the title will be displayed in.

Size is font size in points.

Style choices vary depending on the font, but are usually Bold, Normal and Italic.

☑ **Use CSS values**

This is set to use the values of the workspace presentation cascading style sheet (`ws_pres.css`) for the label text (header, footer, and chart legend). This option is useful for text in double-byte character set (DBCS) languages that are more legible with a different font styling than the defaults. The Tivoli Enterprise Portal administrator can edit the `ws_pres.css` style sheet in the `Install_dir`\cnb (Windows); cw (Linux)\classes\candle\fw\ resources\styles directory on the computer where the Tivoli Enterprise Portal Server is installed, to establish a specific font style that will be used for all label text and that will not be overwritten whenever your installation is upgraded to a new release.

Clear the check box if you want to use the font styling that is specified in the **Text style** area.

☑ **Always connects points**

This option is only displayed for Plot Charts and is enabled by default. The plot points are connected with lines even when gaps in time are present.

Advanced Filters

The Filters tab of the Properties editor has a 🔀 **Set advanced filters** button that opens this window, where you can specify a sort order for the returned values and the highest or lowest values to return.

☐ **Enable First/Last Filter**

Select the check box to filter the view to show the first or last rows returned based on the value of the attribute specified in the **Sort by** area.

A first/last filter is different from the **First/Last Functions** in the Advanced Query window because it applies the filter after the data has been returned (rather than as part of the query) so that all the rows that were returned are reviewed for inclusion in the view. For example, instead of the last 10 rows from each remote monitoring server in the view, which would display a total of 50 rows from five monitoring servers, the last 10 rows from all monitoring servers are shown.

Sort by

Specify a column (from the 🔽 list) to sort by when selecting the first or last rows and in which direction. This sort has to do only with sorting of the returned rows

so that the first or last rows can be determined. It has no affect on any other sort that you might have specified for the view.

◉ **Ascending**

An ascending sort causes the attribute selection to be from the first rows returned.

◯ **Descending**

A descending sort causes the attribute selection to be from the last rows returned.

First Last

Specify the number of rows to return and whether they should be from the first rows of the returned list or from the last rows.

◉ **First**

This causes the attribute selection to be from the first rows returned.

◯ **Last**

This causes the attribute selection to be from the last rows returned.

Example

This example is a table view that shows the results of selecting the top 5 processes in descending order, based on the Working Set attribute: Sort by ▾ Working Set ◉ **Descending**; ◉ **First** 5.

Process Name	Virtual Bytes	Page File Bytes	Working Set
java	411,787,264	107,171,840	*132,677,632*
javaw	669,335,552	56,926,208	*37,728,256*
svchost	137,080,832	15,708,160	*15,097,856*
kntcma	103,129,088	12,484,608	*9,277,440*
winlogon	63,881,216	8,871,936	*7,200,768*

The original results of a custom Windows Process query are shown for comparison.

Process Name	Virtual Bytes	Page File Bytes	Working Set
smss	3,895,296	184,320	245,760
csrss	29,405,184	2,105,344	2,318,336
winlogon	63,881,216	8,871,936	7,200,768
services	37,478,400	2,117,632	1,945,600
lsass	42,614,784	3,899,392	2,277,376
svchost	64,270,336	2,916,352	2,285,568
svchost	137,080,832	15,708,160	15,097,856
svchost	30,408,704	1,323,008	1,392,640
svchost	37,830,656	1,708,032	1,642,496
spoolsv	44,396,544	3,563,520	2,973,472
jservice	621,842,432	16,556,032	4,227,072
ccEvtMgr	42,061,824	2,551,808	438,272

Process Name	Virtual Bytes	Page File Bytes	Working Set
ccSetMgr	33,665,024	2,551,818	704,512
c4ebreg	43,421,696	6,254,592	2,592,768
issimsvc	37,322,752	5,443,584	1,708,032
SavRoam	35,893,248	1,404,928	1,028,096
javaw	669,335,552	56,926,208	37,728,256
java	411,787,264	107,171,840	132,677,632
khdxprto	94,371,840	11,431,936	5,025,792
kntcma	103,129,088	12,484,608	9,277,440
kuma620	91,131,904	9,035,776	3,305,472
kinconfg	49,127,424	1,859,584	2,908,160
logonui	34,713,600	2,875,392	3,330,048

Table view

The table view is a query-based workspace view of data samples collected for an attribute group, or from an ODBC or JDBC data source when a custom SQL query has been written.

The ⊞ table view shows a column for each attribute and either a single row of data (for single-row attribute groups) or multiple rows (for multiple-row attribute groups or historical reporting). For example, the NT Process attribute group displays a row for every process, and NT Memory displays a single row of memory statistics for the agent's host.

If historical data collection is configured and started for the attributes used by the query that retrieves data for the view, you can also view historical data by setting the time span for data to be viewed using the ⬚ Time span tool.

Adding a table view

Use the table view to display data in the workspace. Data samplings from the attributes specified in the query are used to populate the view.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

About this task

Take these steps to add a table view to a workspace:

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new space, click ▯ **Split Vertically** or ▭ **Split Horizontally** in an existing view. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty. If you want to replace a view with a different type of view, skip this step.
3. Click ⊞ **Table**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.

4. Click inside the view where you want the table. If no data fills the table, you must choose a query for it.
5. If a message asks you to select a query now, click **Yes** to open the Properties editor. Specify the data to include and the formatting of the table view:

Option	Description
Query	Select a query for the data to retrieve into the view. See "Query" on page 92.
Filters	Specify the attributes whose values you want to exclude or include in the table and, if you like, show only those values that fall within a range. See "Filters" on page 93.
Thresholds	Highlight cells that fall outside a threshold value. See "Thresholds" on page 95.
Style	Change and enhance the appearance of the table. See "Style" on page 96.

Results

The query collects data samplings from the managed systems that are assigned to the Navigator item and populates the view with them.

What to do next

To see the assigned managed systems, right-click the current Navigator item and click Properties.

If the attribute group specified in the query has historical data collection enabled, you can specify a time period to show.

To keep the table view in this workspace for future work sessions, either click **Save** to update the workspace with the new view or click **File** > **Save Workspace As** to save this as a new workspace and leave the original workspace unchanged.

Table view features

The table view has features for refining the display to suit your needs. The table view features allow you to refine the display of your data. These features also apply to the message log and event console views.

Column descriptions
Descriptions of the attributes for which data is reported are displayed as hover help when you move the mouse pointer over a column heading.

Column adjustments
You can make the following adjustments to columns:

Adjust column width by dragging a border left or right.

Lock a column by right-clicking in the column heading and selecting Lock this Column. The column and any columns to left are fixed in place as you scroll the table left or right. To unlock, right-click a column heading and select Unlock Table Columns.

Change column order by dragging a column heading and dropping to insert it between two columns.

Column sort by clicking a column heading to sort the table by that column in ascending order, clicking again for descending, and again to go back to the original arrangement. You can also build a query that specifies a sort order, which is necessary for multiple-page tables if you want all the rows sorted and not just those in the current page.

Find

The Find function enables you to enter search criteria for the columns shown in the table.

Threshold indicators

Threshold cells have either a severity icon or a colored background to indicate values that meet the threshold criteria. A threshold column heading has an indicator, such as ⚠ Minor, that you can click to sort the column by the threshold value. If the column has multiple thresholds assigned, the highest level indicator is displayed. When threshold cells do not show in the viewable area, the scroll button is colored or has a severity icon to show where to click to see them.

Link indicators

🔗 Link indicators appear to the left of a row if the table has a link filtered by one or more row values. Click the 🔗 indicator to open the default linked workspace or, if no default link has been assigned, a list of available links. You can also right-click the indicator to get a list of available links or to open the link wizard to define a link.

Multiple pages

Multiple pages, a page being 100 rows, are indicated by ⏫ and ⏬ buttons on the scroll bar. A text box is displayed at the top of the table with "Page 1 of ". You can type another number to go to that page. Replace 1 with a high number such as 99 and press **Enter** to go to the last page and adjust the total number of pages displayed to the correct number.

Click ✎ Edit Properties in the view title bar to open the Properties editor, and change the **Page size** in the Query tab. A smaller page size speeds up data retrieval because only one page is retrieved at one time rather than the entire row set.

Historical reporting

As well as seeing real-time data in the table view, you can have Tivoli Enterprise Portal log the data into history files so you can examine data over longer periods of time. The table view has a 🕑 **Time Span** tool for setting a time period to report. The table retains data from previous refreshes and appends newly arrived data to the view.

Adding column thresholds

You can decorate a cell in the table view (including event console, message log, and universal message views) with a background color, severity icon, or custom icon when the value exceeds a given threshold.

About this task

You can decorate a cell with a background color, severity icon, or custom icon by writing formula expressions that specify the threshold criteria and the indicators to use, such as `1908 ⊙` or `1908` .

1. If the Properties editor is not open, click ✎ Edit Properties in the table view title bar.
2. Click the ▣ **Thresholds** tab.

3. If you want to place a severity icon in a threshold cell instead of a background color, select the **Use Icons** ☑ check box.

4. Click inside a cell in the Threshold ⊗ column (second column) to see and select from the list of severities.

   The seven predefined severities, from Fatal to Unknown, display in the list. If you are selecting an icon, also shown are any custom severity icons that were added to the portal server. The icons are ordered by their severity, from highest to lowest.

5. In the same row as the selected highlight, click in a column cell and write the threshold expression for that attribute. This activates the editor and formula bar:

   Click ▼ if you want to change the function from the default **Value of expression** (or **Compare date/time** for time attributes) to another function.

   Click == if you want to change the relational operator from Equal to Not equal or to any other operator in the list.

   Click inside the text box to enter the threshold value. Enumerated attribute values, such as *online* or *offline*, have a ▼ list of possible choices.

6. Click **Test** to evaluate the threshold expression and to see in the Preview area how the results will be displayed.

7. Repeat steps 3 to 5 to add more threshold expressions to this or another column:

   The threshold expressions entered in the same row have an AND relationship. When a retrieved value satisfies all the expressions in a row, the cell will be highlighted.

   The threshold expressions in each row are independent; they have an OR relationship. If a retrieved value satisfies the conditions in more than one row, the cell will be highlighted with the highest severity.

8. When the threshold formula is complete, save your changes by clicking **Apply** if you want to leave the Properties editor open or **OK** to close the editor.

Note:

1. Values in the table that meet the threshold criteria are indicated by the severity background color or severity icon. You will notice a severity icon in the column heading of threshold cells. This is the highest threshold severity established for the column; click the icon to sort the view by severity.

2. The scroll buttons indicate threshold cells that are outside the visible area. For example, ▶◆ guides you to scroll right to see more informational threshold cells. Move the mouse pointer over a threshold cell to see the threshold criteria.

3. In the event that multiple threshold expressions written for the same column become true at the same time, the severity with the highest value will be applied. If the same severity was applied (which is possible only if the icon was selected for one expression and its background color counterpart for the other), then the first row expression will determine the appearance. For example, if the icon was specified in row 2 of the formula editor and the background color in row 3, then the icon will be used.

Guidelines for threshold icons

There are already seven predefined severity icons that can be assigned to threshold cells, and you can add your own severity threshold icons to the portal server.

When adding your own severity threshold icons to the portal server, observe these requirements and guidelines:

- Use iconic images sized at 9 pixels by 9 pixels to match the size of the predefined severity icons. You can use larger images, but they will be clipped when applied to a scroll button.
- Save the images in the GIF format. Nonlinear GIFs look best when saved with a transparent background.
- Include the severity value in the filename, immediately preceding the .gif extension (such as icon28.gif, serious_92.gif, or 999.gif). Valid values are greater than 0 and less than MAX_INT and must be unique. If a duplicate severity is found, the icon file will not be included in the list of possible icons that can be used for threshold cells. These are the severity values of the predefined icons:
  - ⊗ Fatal = 100
  - ⊗ Critical = 90
  - ⚠ Minor = 40
  - ⚠ Warning = 30
  - ☑ Harmless = 27
  - ▣ Informational = 25
  - ◉ Unknown = 15
- Use only ASCII characters in the filename. Names that use characters outside the ASCII code range will not be displayed in the list of available files when the user browses to select one.
- On the computer where the Tivoli Enterprise Portal Server installed, place the GIF files in the user directory:

  Linux or **UNIX** `<itm_installdir>/cw/classes/candle/fw/resources/icons/user`

  Windows `<itm_installdir>\cnb\classes\candle\fw\resources\icons\user`
- If users are logged on to the portal server, they will need to exit the portal and log on again before the icons will be available for inclusion in the threshold criteria.

Chart views

Add one of the seven available chart views to a workspace to present data for managed systems in real time. When a query is assigned to a chart for an attribute group that has historical data collected, you can also display historical data.

Each type of chart is best suited to displaying a certain type of data:

Chart type	Best for...	Multiple attributes	Multiple rows	Multiple managed systems	🕒 Time span
◔	Showing proportional value to the whole	One slice per attribute	One pie for each row	One pie for each managed system	Yes
▥ [1]	Comparing values among related attributes	One bar per attribute. Stacking bars show one segment per attribute	One set of bars for each row. Stacking bars show one bar per row	One set of bars for each managed system	Yes

Chart type	Best for...	Multiple attributes	Multiple rows	Multiple managed systems	▦ Time span
▦ [1]	Showing trends over time and among related attributes	One line per attribute; one data point for each data sampling	Yes	Yes	Yes
▦ [1]	Showing trends over time and among related attributes	One line per attribute; one data point for each data sampling	Yes	Yes	Yes
◕	Showing individual elements that change frequently	One gauge per attribute	No	No	No
▯	Showing cumulative values	One gauge per attribute	No	No	No
[1] This chart can also be used to apply a baseline for predictive analysis. For more information, see "Chart baselines" on page 126.					

Creating a new chart

Add a chart view to a workspace to better see trends and data spikes in sampled data. The chart types available are ◕ Pie, ▥ Bar, ▦ Plot, ▦ Area, ◕ Circular gauge, and ▯ Linear gauge.

Before you begin

Consider the type of data that you are creating a chart for and select the chart type that best illustrates the data. For example, you might want a gauge chart to show individual attribute values that fluctuate frequently, or you might want a pie chart to show proportional values.

Your user ID must have Workspace Author Mode permission to use this function.

About this task

Take these steps to create a chart view:

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new space, click ▥ **Split Vertically** or ▤ **Split Horizontally** in an existing view. If you do not see the ▥ or ▤ tools, the Split Pane Toolbar has been disabled. Select it from the View menu to enable it. You can also right-click inside a view and select ▥ **Split Vertically** or ▤ **Split Horizontally**. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
3. Click the tool for the type of chart you want:

Option	Description
	Select the pie chart to show proportional values to the whole, one slice per attribute, one pie for each row returned, and one pie for each managed system if the Navigator item has multiple managed systems assigned.
	Select the bar chart for comparing values among related attributes, one bar per attribute or per segment if this is a stacking bar chart, one set of bars for each row, and one set of bars for each managed system if the Navigator item has multiple managed systems assigned. You can apply a monitored baseline of a situation's threshold criteria to a bar chart to see how effective the situation is. See "Adding monitored baselines to a chart" on page 126.
	Select the plot chart for showing trends over time and among related attributes, one line per attribute and one data point for each data sampling, with multiple rows and managed systems supported within a single chart. This chart can be used to perform any or all of the three baselining functions for analysis. See "Chart baselines" on page 126.
	Select the area chart for showing trends over time and among related attributes, one line per attribute and one data point for each data sampling, with multiple rows and managed systems supported within a single chart. You can apply a monitored baseline of a situation's threshold criteria to a bar chart to see how effective the situation is, or apply statistical functions to a plotted attribute. See "Adding monitored baselines to a chart" on page 126 and "Adding statistical baselines to a chart" on page 130.
	Select the circular gauge chart for showing individual elements that change frequently, with one gauge per attribute displayed. Neither multiple rows nor multiple managed systems are supported for a single circular gauge.
	Select the linear gauge chart for showing cumulative values. Neither multiple rows nor multiple managed systems are supported for a single linear gauge.

The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.

4. Click inside the space to adopt the new chart. Either the Select Attribute window opens or a message asks if you want to assign a query.

5. If the Select Attribute window opens, select an attribute (use Ctrl+Click or Shift+Click to select more), then click **OK**. This happens when the original view already has a query assigned. If this is not the query that you want for the chart, select one of the numeric attributes to complete the operation, then continue to the next step.

6. If you are prompted to assign the query now, click **Yes**, then choose a query and select the attributes to include:

   a. In the **Query** tab, click [icon] **Click here to assign a query** to open the Query editor. The Query editor shows a tree on the left organized by Tivoli monitoring product and then by attribute group.

   b. Click [icon] to expand the [icon] Monitoring agents and [icon] attribute group folders to see the queries.

   c. Select a [icon] Query title to see its description. The query description, edit date, and specification are displayed in the right frame. The selected options indicate which attributes from the group are included in the query. The cells in the numbered rows show any filter criteria for the query.

   d. Click **OK** to select the query and return to the Properties editor. The Description and Formula areas of the Query tab show the name, description, and formula of the query that you selected; the Preview area shows how the chart will look when the query is applied.

   e. Click the **Filters** tab, select the check box under a column heading to include its value in the chart. Charts can show only numeric values (no text or timestamp attributes); a disabled [icon] check box means that the attribute cannot be plotted.

   As you edit the properties, click **Test** to see your changes in the Preview area at the top of the window.

7. Click **Apply** to save your changes and keep the Properties editor open, or **OK** to save your changes and close the editor.

8. If you want to keep the view in this workspace for future work sessions, save the workspace or keep the original intact and create a new workspace:

   • Click [icon] **Save** to update the workspace properties with the new view.

   • Click **File** > **Save Workspace As** to save this as a new workspace and leave the original workspace unchanged since the last time it was saved.

Results

Data samples from the attributes specified in the query and the Filters tab are plotted on the chart. If no information appears in a table or chart, it might be because no data samplings meet the query or filter criteria.

If you created a gauge chart, it can show data from only one managed system regardless of the managed systems assigned to the Navigator item. If you are using the query in a workspace where multiple managed systems are applicable (at the platform level of the Navigator Physical view), the data shown is from the first managed system in the Assigned list.

What to do next

See the topic for the chart type you have chosen for a description of the chart's unique aspects and options, with links to styling options that are shared with other chart types.

If you have historical collection configured for the queried attribute group, you can also display data from previous refreshes using ▧ **Time Span** in the view toolbar. (Not available for gauge charts)

Tip: If you want to use the same query for two charts, create the first view, then split the view to adopt the same query in the other chart, then apply a different chart type to the new view.

Related reference:

"Pie chart"
Add a pie chart to a workspace to show the proportional value of each attribute to the whole. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Bar chart" on page 109
Add a bar chart to a workspace to compare values among related attributes. The stacking bar chart best shows multiple values for the same attribute. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Plot chart" on page 110
Add a plot chart to a workspace to show trends over time among related attributes. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Area chart" on page 112
Use the area chart for showing trends over time among related attributes. The area chart is like the plot chart except that the area below the plotted line is filled in with color to indicate volume.

"Circular gauge chart" on page 114
Add a circular gauge chart to a workspace to show individual elements that change frequently. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options.

"Linear gauge chart" on page 115
Add a linear gauge chart to a workspace to show individual elements that change frequently. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options.

"Chart style" on page 117
The 𝐴 **Style** tab in the Properties editor to refine the appearance of the chart. For the bar, plot, and area charts, this is also where you can specify a line chart overlay and threshold indicators.

Pie chart

Add a pie chart to a workspace to show the proportional value of each attribute to the whole. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

A pie chart has a slice for every data point in a single data series (row). If you specify attributes from a group that returns multiple rows, a pie is drawn for each row. Similarly, when the query is assigned to multiple managed systems, you get a pie for each managed system.

You can tilt a pie chart by holding down the Alt key while dragging the mouse pointer up or down anywhere on its surface; adjust the height of the pie by holding down the Shift key while dragging the mouse pointer.

After adding the chart, use the styling options in the Properties editor to refine the appearance. If the Properties editor is not open, click ✎ **Edit properties** in the view's title bar. The A꜀ **Style** tab shows a thumbnail graphic of the chart with areas that you click to edit the corresponding options.
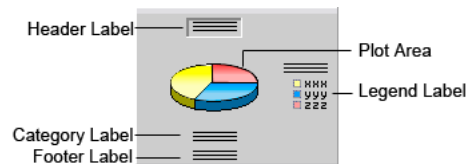
Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor A꜀ **Style** tab to add a title to the view, which shows in the view title bar. You can also apply the title as a view header, add a border, and change the text style. The default name is **Pie Chart**. See "Header Label" on page 117.

Plot Area

Select the pie in the thumbnail graphic to open the Pie Chart tab with these options:

Effect The pie chart has a three-dimensional appearance by default, and you can flatten the appearance by switching to 2D.
- ○ **2D** to render a two-dimensional chart.
- ◉ **3D** to render a three-dimensional chart.

Pie Slice Offsets
You can specify that certain pie slices should be offset from the rest of the chart and labeled for emphasis.
- In the **Offset Value** field, specify the offset distance in pixels.
- Specify the pie slices you want to offset by selecting the appropriate check boxes.
- Click the icon beside the name of a pie slice to open the Chart Label options window, from which you can control options affecting how the offset slices are labeled.

Legend Label

Select the Legend Label area of the thumbnail graphic in the Properties editor A꜀ **Style** tab to control the display of the chart legend and its appearance, including the coloring of the data series: pie slices, bars, plot lines, and plot areas. See "Legend Label" on page 123.

Category Label

Select the Category Label area of the thumbnail graphic to determine whether to display a label for the data series, one for each pie in the chart, and the attribute to use for the label.

Text The **Text** area determines the font to use, its size and whether any styling is applied. Select any changes from the ▾ list for **Font**, **Size**, or **Style**.

Attribute

You can either accept the Default Category Label for the chart type or select an attribute from the ▾ list.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor ᴀ **Style** tab to add a label at the bottom of the view. See "Footer Label" on page 126.

Related tasks:

"Creating a new chart" on page 104
Add a chart view to a workspace to better see trends and data spikes in sampled data. The chart types available are ⊙ Pie, ▥ Bar, ▨ Plot, ▨ Area, ⊕ Circular gauge, and ▮ Linear gauge.

"Setting a time span to display" on page 44
You can have the Tivoli Enterprise Portal log data samplings into history files or a database for display in a table or chart. These historical data enabled views have a tool for setting a range of previous data samples to be reported.

Bar chart

Add a bar chart to a workspace to compare values among related attributes. The stacking bar chart best shows multiple values for the same attribute. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

In the bar chart, the attributes are listed vertically with each attribute given its own bar, and values drawn horizontally. If you specify multiple-row attributes, a bar grouping is added for every row.

You can rotate a 3D bar chart by holding down the Alt key while dragging the mouse pointer anywhere on the chart; adjust the thickness of a bar by holding down the Shift key while dragging the mouse.

If the bar chart has limited space for showing category labels, you see labels for every other bar, every third bar, or whatever is required to fit the space. Try ◻ maximizing the view to see all the labels and then ⧉ restore when you are done.

After adding the chart, use the styling options in the Properties editor to refine the appearance. If the Properties editor is not open, click ✎ **Edit properties** in the view's title bar. The ᴀ **Style** tab shows a thumbnail graphic of the chart with areas that you click to edit the corresponding options.

Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor ᴀ **Style** tab to add a title to the view, which shows in the view title bar. You can also apply the title as a view header, add a border, and change the text style. The

default name is **Pie Chart**. See "Header Label" on page 117.
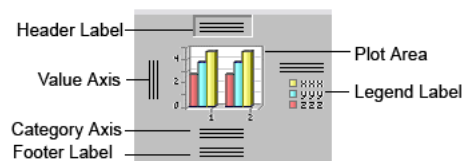
Plot Area

Click the bars in the thumbnail graphic to open the plot area tabs: Bar Chart, Chart Thresholds, and Chart Markers. See "Plot Area" on page 117.

Legend Label

Select the Legend Label area of the thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to control the display of the chart legend and its appearance, including the coloring of the data series: pie slices, bars, plot lines, and plot areas. See "Legend Label" on page 123.

Value and Category Axes

Edit the axis style properties to switch between vertical and horizontal orientation, and to add labels and styling options. See "Value and Category Axes" on page 125.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to add a label at the bottom of the view. See "Footer Label" on page 126.

Related tasks:

"Creating a new chart" on page 104
Add a chart view to a workspace to better see trends and data spikes in sampled data. The chart types available are Pie, Bar, Plot, Area, Circular gauge, and Linear gauge.

"Setting a time span to display" on page 44
You can have the Tivoli Enterprise Portal log data samplings into history files or a database for display in a table or chart. These historical data enabled views have a tool for setting a range of previous data samples to be reported.

"Adding monitored baselines to a chart" on page 126
Apply a monitored baseline of a situation's threshold criteria to a bar chart, plot chart, or area chart to see how effective the situation is in this environment. You can visually compare the current values in the chart with the situation and adjust the thresholds as needed.

Plot chart

Add a plot chart to a workspace to show trends over time among related attributes. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

A plot chart shows changes over a period of time by drawing a continuous line from one data point to the next, one data point for each data sampling and one line for each attribute selected. This type of chart is useful for showing trends over time and among related attributes.

Each time you open the workspace or edit its properties, the plot chart starts anew. The first set of data points is plotted, the next set is plotted after the refresh interval, and a line is drawn from each previous plot point to the current one. If historical data is being collected and a time span has been specified, the plot chart

is populated with the historical data (unless data collection was started only recently and historical data has not been gathered yet for the period specified).

The redistribution and restarting of data collection can result in some gaps in the data if the collection interval is frequent, because collection is aligned with time boundaries. Time gaps in your data might affect your plotted points. In some cases lines might not be drawn to connect plotted points.

After adding the chart, use the styling options in the Properties editor to refine the appearance. If the Properties editor is not open, click ✐ **Edit properties** in the view's title bar. The 𝐴ₐ **Style** tab shows a thumbnail graphic of the chart with areas that you click to edit the corresponding options.

Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to add a title to the view, which shows in the view title bar. You can also apply the title as a view header, add a border, and change the text style. The default name is **Pie Chart**. See "Header Label" on page 117.

Plot Area
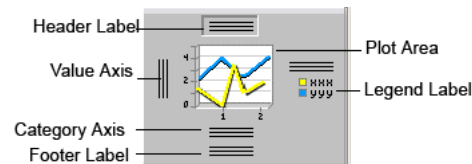
Click the plot area in the thumbnail graphic to open these tabs: Plot Chart, Chart Thresholds, and Chart Markers. See "Plot Area" on page 117.

Legend Label

Select the Legend Label area of the thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to control the display of the chart legend and its appearance, including the coloring of the data series: pie slices, bars, plot lines, and plot areas. See "Legend Label" on page 123.

Value and Category Axes

Edit the axis style properties to switch between vertical and horizontal orientation, and to add labels and styling options. See "Value and Category Axes" on page 125.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to add a label at the bottom of the view. See "Footer Label" on page 126.

Related tasks:

"Creating a new chart" on page 104
Add a chart view to a workspace to better see trends and data spikes in sampled data. The chart types available are ⊕ Pie, 📊 Bar, ▦ Plot, ▦ Area, ⊕ Circular gauge, and 📊 Linear gauge.

"Setting a time span to display" on page 44
You can have the Tivoli Enterprise Portal log data samplings into history files or a database for display in a table or chart. These historical data enabled views have a tool for setting a range of previous data samples to be reported.

"Adding monitored baselines to a chart" on page 126
Apply a monitored baseline of a situation's threshold criteria to a bar chart, plot chart, or area chart to see how effective the situation is in this environment. You can visually compare the current values in the chart with the situation and adjust the thresholds as needed.

"Adding statistical baselines to a chart" on page 130
Use the statistical baseline function on a plot chart or area chart to apply any of these statistical functions to the selected attribute: Range between minimum and maximum, average, minimum, maximum, percentile, and mode.

"Adding historical baselines to a chart" on page 132
Use the values from the current plot chart to identify daily trends in threshold conditions. You can then better anticipate problems by comparing current values with expected future values or with past values, and overlaying the chart with situation comparison criteria, one plotted line for each day specified.

Area chart

Use the area chart for showing trends over time among related attributes. The area chart is like the plot chart except that the area below the plotted line is filled in with color to indicate volume.

When multiple attributes are included, the first attribute is plotted as a line with color fill followed by the second attribute, and so on. A higher values obscures a lower value if the lower value was plotted first. If this happens, you can rearrange the attributes by dragging the column headings in the Filters tab to arrange them in descending order of value. You can also switch to a stacked area.

After adding the chart, use the styling options in the Properties editor to refine the appearance. If the Properties editor is not open, click ✎ **Edit properties** in the view's title bar. The 🄰 **Style** tab shows a thumbnail graphic of the chart with areas that you click to edit the corresponding options.

Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor 🄰 **Style** tab to add a title to the view, which shows in the view title bar. You can also apply the title as a view header, add a border, and change the text style. The default name is **Pie Chart**. See "Header Label" on page 117.

Plot Area
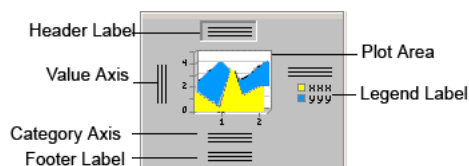
Click the plot area in the thumbnail graphic to open these tabs: Area Chart, Chart Thresholds, and Chart Markers. See "Plot Area" on page 117.

The following screen captures are the same area chart showing ◼ % Processor Time, ◼ % Privileged Time, and ◼ % User Time. **1** is the initial rendering. The % Privileged Time value is almost completely obscured by the higher % User Time value. **2** is the chart after moving the % Privileged Time column heading after % User Time in the Properties editor **Filters** tab so that the highest value is plotted first, followed by the next highest value, and then the lowest value. Now all the values are revealed. **3** is the chart after changing **Type** to "Stacking Area Chart", whereby the values are stacked one upon the other.

Legend Label

Select the Legend Label area of the thumbnail graphic in the Properties editor *A₄* **Style** tab to control the display of the chart legend and its appearance, including the coloring of the data series: pie slices, bars, plot lines, and plot areas. See "Legend Label" on page 123.

Value and Category Axes

Edit the axis style properties to switch between vertical and horizontal orientation, and to add labels and styling options. See "Value and Category Axes" on page 125.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor *A₄* **Style** tab to add a label at the bottom of the view. See "Footer Label" on page 126.

Related tasks:

"Creating a new chart" on page 104
Add a chart view to a workspace to better see trends and data spikes in sampled data. The chart types available are Pie, Bar, Plot, Area, Circular gauge, and Linear gauge.

"Setting a time span to display" on page 44
You can have the Tivoli Enterprise Portal log data samplings into history files or a database for display in a table or chart. These historical data enabled views have a tool for setting a range of previous data samples to be reported.

"Adding monitored baselines to a chart" on page 126
Apply a monitored baseline of a situation's threshold criteria to a bar chart, plot chart, or area chart to see how effective the situation is in this environment. You can visually compare the current values in the chart with the situation and adjust the thresholds as needed.

"Adding statistical baselines to a chart" on page 130
Use the statistical baseline function on a plot chart or area chart to apply any of these statistical functions to the selected attribute: Range between minimum and maximum, average, minimum, maximum, percentile, and mode.

Circular gauge chart

Add a circular gauge chart to a workspace to show individual elements that change frequently. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options.

The circular gauge chart shows the proportional amount of a data series, one gauge for each attribute chosen. This type of chart is well suited for showing individual elements that change frequently, such as percentage user time.

If the attribute displayed is from a multiple-row attribute group, only the value from the first row is charted. The circular gauge chart can show data from only one managed system.

After adding the chart, use the styling options in the Properties editor to refine the appearance. If the Properties editor is not open, click **Edit properties** in the view's title bar. The **Style** tab shows a thumbnail graphic of the chart with areas that you click to edit the corresponding options.
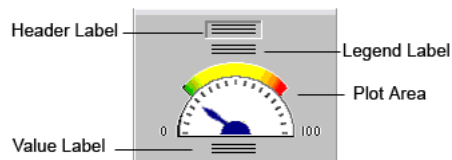
Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor **Style** tab to add a title to the view, which shows in the view title bar. You can also apply the title as a view header, add a border, and change the text style. The default name is **Pie Chart**. See "Header Label" on page 117.

Plot Area

Click the circular gauge in the thumbnail graphic to open these tabs: Scale, Center, Needle, Tick, and Range. See "Plot Area" on page 117.

Legend Label

Use this tab to specify the colors and text associated with the data points and data series in the chart.

Legends
Lists the attributes that are being plotted. Select an attribute from the list to see and change the legend text or the plot point (or series) color.

Text   The legend label text. By default, this is the attribute name.

Value Label

Select the Value Label area of the gauge chart thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to show or hide the attribute's value at the bottom of the view or to edit the font.

☐ **Show**
Select the check box to display the value of the attribute at the bottom of the view.

Text style
Select a different **Font name** or font **Size** from the ▾ list. ☑ **Bold** is used to boldface the value. ☐ **Italic** is used to italicize the value.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor 𝐴ₐ **Style** tab to add a label at the bottom of the view. See "Footer Label" on page 126.

Related tasks:

"Creating a new chart" on page 104
Add a chart view to a workspace to better see trends and data spikes in sampled data. The chart types available are ⬤ Pie, ⬛ Bar, ⬛ Plot, ⬛ Area, ⬤ Circular gauge, and ⬛ Linear gauge.

Related reference:

"Chart style" on page 117
The 𝐴ₐ **Style** tab in the Properties editor to refine the appearance of the chart. For the bar, plot, and area charts, this is also where you can specify a line chart overlay and threshold indicators.

Linear gauge chart

Add a linear gauge chart to a workspace to show individual elements that change frequently. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options.

The linear gauge chart shows the collective value of every item in a single data series, one gauge for each attribute chosen.

If the attribute displayed is from a multiple-row attribute group, only the value from the first row is charted. The linear gauge chart can show data from only one managed system.

After adding the chart, use the styling options in the Properties editor to refine the appearance. If the Properties editor is not open, click ✎ **Edit properties** in the view's title bar. The ᴬᵃ **Style** tab shows a thumbnail graphic of the chart with areas that you click to edit the corresponding options.
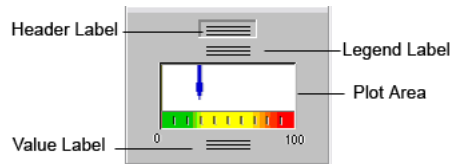
Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor ᴬᵃ **Style** tab to add a title to the view, which shows in the view title bar. You can also apply the title as a view header, add a border, and change the text style. The default name is **Pie Chart**. See "Header Label" on page 117.

Plot Area

Click the linear gauge in the thumbnail graphic to open these tabs: Scale, Needle, Tick, and Range. See "Plot Area" on page 117.

Legend Label

Use this tab to specify the colors and text associated with the data points and data series in the chart.

Legends
> Lists the attributes that are being plotted. Select an attribute from the list to see and change the legend text or the plot point (or series) color.

Text The legend label text. By default, this is the attribute name.

Value Label

Select the Value Label area of the gauge chart thumbnail graphic in the Properties editor ᴬᵃ **Style** tab to show or hide the attribute's value at the bottom of the view or to edit the font.

☐ **Show**
> Select the check box to display the value of the attribute at the bottom of the view.

Text style
> Select a different **Font name** or font **Size** from the ▾ list. ☑ **Bold** is used to boldface the value. ☐ **Italic** is used to italicize the value.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor ᴬᵃ **Style** tab to add a label at the bottom of the view. See "Footer Label" on page 126.

Related tasks:

"Creating a new chart" on page 104
Add a chart view to a workspace to better see trends and data spikes in sampled
data. The chart types available are ⬤ Pie, ▥ Bar, ▦ Plot, ▨ Area, ⬤ Circular
gauge, and ▤ Linear gauge.

Related reference:

"Chart style"
The 𝐴ₐ **Style** tab in the Properties editor to refine the appearance of the chart. For
the bar, plot, and area charts, this is also where you can specify a line chart overlay
and threshold indicators.

Chart style

The 𝐴ₐ **Style** tab in the Properties editor to refine the appearance of the chart. For
the bar, plot, and area charts, this is also where you can specify a line chart overlay
and threshold indicators.

Header Label

Select the Header Label area of the thumbnail graphic in the Properties editor
𝐴ₐ **Style** tab to add a title to the view, which shows in the view title bar. You can
also apply the title as a view header, add a border, and change the text style.

☐ **Show**

Select the check box to display the title as a header.

☐ **Show Border**

Select the check box to display a border around the header label.

Text   This text you enter here becomes the view title and is displayed in the title
bar and, if you selected ☑ **Show**, in the header.

Text style

Select a different **Font name** or font **Size** from the ▼ list. ☑ **Bold** makes
the header text boldface. ☐ **Italic** makes the header text italicized. When ☑
Use CSS values is selected, this area is disabled.

☑ **Use CSS values**

This is set to use the values of the workspace presentation cascading style
sheet (`ws_pres.css`) for the label text (header, footer, and chart legend).
This option is useful for text in double-byte character set (DBCS) languages
that are more legible with a different font styling than the defaults. The
Tivoli Enterprise Portal administrator can edit the `ws_pres.css` style sheet
in the *Install_dir*`\cnb (Windows); cw (Linux)\classes\candle\fw\`
`resources\styles` directory on the computer where the Tivoli Enterprise
Portal Server is installed, to establish a specific font style that will be used
for all label text and that will not be overwritten whenever your
installation is upgraded to a new release.

Clear the check box if you want to use the font styling that is specified in
the **Text style** area.

Plot Area

The chart data is rendered in the plot area. The plot area has such styling options
as stacking the data series and adding threshold indicators.

The tabs that are displayed depend on the chart type: "<Chart type> Chart" on
page 118 for pie charts, bar charts, plot charts and area charts; "Chart Thresholds"
on page 119 and "Chart Markers" on page 119 for bar charts, plot charts, and area

charts; and "Scale" on page 120, "Center" on page 121, "Needle" on page 121, "Tick" on page 121, and "Range" on page 121 for gauge charts.

<Chart type> Chart

Pie charts, bar charts, plot charts and area charts have this tab for effects that are unique to the chart.

Type For bar charts and area charts, you can change the chart to a stacked chart, which shows cumulative values. The Stacking Bar Chart is useful for showing multiple values for the same (multi-row) attribute.

Effect For bar charts and pie charts, the appearance is three-dimensional by default and you can flatten the appearance by switching to 2D:

○ **2D** to render a two-dimensional chart.

◉ **3D** to render a three-dimensional chart.

Select Plot Points

For plot charts, select ◉ **Attributes in first row** if you are plotting a single-row attribute group with the current data samples; or ○ **Attributes across multiple rows** if this is a multiple-row attribute group or you are plotting historical data.

Plot Duration

For plot charts and area charts, the chart duration is set to one hour initially. After that, as new data samples arrive, the oldest plot points are removed to make room for the new. To change the duration from as little as five minutes to as much as 24 hours, either drag the slider bar or enter the time in the **Hours** and **Minutes** text boxes.

☐ **Included in axis bounds**, when selected, takes into account the elapsed time when calculating the axis end time. The effect is that the chart rendering stops where the last data sample was plotted rather than stretch across the entire plot area. After the chart duration has been exceeded both types of charts look the same.The Chart Duration area is disabled if this is a historical chart.

Refresh Rate

For plot charts and area charts, the refresh rate determines how often a data sample is taken at the agent and plotted on the chart. ◉ **Use Refresh Rate from Workspace** is the default selection. The refresh rate that is set for the workspace (View → Refresh Every) is applied to this chart as well. Select ○ **Refresh Area Chart Every** to specify a refresh rate for the chart that is independent of the workspace refresh rate. You can choose a value from 10 seconds to 60 minutes.

Overlay

For bar charts, plot charts, and area charts, you can click **Add Overlay** to open a window for overlaying the chart with a line chart. Here you can add other attributes from the view's query that are applied as plot points and connecting lines. See "Add Overlay" on page 122.

Here is an example of a bar chart with overlay. The chart shows the ■ current size and ■ maximum size of three monitored logs: Application, Security, and System. The ● percentage of the log used is plotted over the bars.
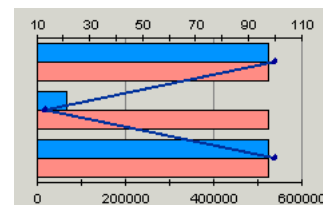
Chart Thresholds

Bar charts, plot charts, and area charts have a Chart Thresholds tab that you can use to define a highlighted region representing a range of values on the chart.

Defined

The list shows the threshold names that are currently defined for the chart. To create a new threshold, click **Create New** and specify the name for the new threshold when prompted. You can then define the properties of the threshold in the **Threshold** and **Fill Style** fields. You can also rename or delete a threshold from the list.

Threshold

The fields in the **Threshold** area of the Chart Thresholds tab define value highlighted by the selected marker, as well as the style of highlighting to use:

☑ **Show Threshold**

The check box is selected to display the threshold on the chart. You can hide a defined threshold by deselecting the check box.

Associated with

Select an available axis to use for thresholding from the ▾ list.

Start value

Accept the lowest value on the chart, -MAX ☑ **Use default**, or deselect the check box and enter the lower boundary of the threshold.

End value

Accept the highest value on the chart, MAX ☑ **Use default**, or deselect the check box and enter the upper boundary of the threshold.

Fill Style

Select the **Pattern** and **Color** to use to highlight the threshold. Click **Custom Colors** to choose from a wider selection of colors or to create your own.

Advanced Options

Click this button to open the Chart Thresholds - Advanced options window, where you can configure additional options for how the threshold is displayed:

Select ☐ **Visible in legend** to include a color indicator and the threshold definition name in the legend.

Select ☐ **Included in data bounds** to confine the threshold display within the data boundary. This field might not change the chart display if the threshold is obscured by data points.

In the **Boundary Lines** area, you can specify that you want to draw lines on the chart marking the start and end values of the threshold, and configure the **Pattern**, **Width**, and **Color** of the boundary lines.

Chart Markers

A marker is a colored or patterned line drawn across the chart at a defined value along the value or category axis. You can use a marker to highlight a particular value for emphasis. Chart markers are available for the bar chart, plot chart, and

area chart views.
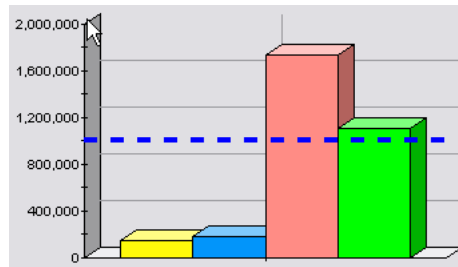
Defined

> The list shows the markers that are currently defined for the chart. To create a new marker, click **Create New** and specify the name for the new marker when prompted. You can then define the properties of the marker. You can also rename or delete a marker from the list.

Marker

> The fields in the **Marker** area of the Chart Markers tab define value highlighted by the selected marker, as well as the style of highlighting to use:
>
> - Click **Show Marker** to display the marker on the chart. You can hide a defined marker by deselecting this check box.
> - In the **Associated with** field, select either the value axis or the category axis.
> - In the **Value** fields, specify the value highlighted by the marker, or select **Use default**. The default value is the lowest value on the chart.
> - In the **Line Style** area, select the pattern, width, and color to use for the marker. Click **Custom Color** to choose from a wider selection of colors or to create your own.

Advanced Options

> Click **Advanced Options** to open the Chart Markers - Advanced options window, where you can configure additional options for how the marker is displayed:
>
> - Select **Visible in legend** to identify the marker by name in the legend for the chart.
> - Select **Included in data bounds** to scale the chart so the marker is shown, even if it is not within the charted data values.
> - Select **Drawn before data** to draw the marker first, so it is displayed behind the charted data. (By default, the marker is displayed in front of the data.)
> - In the **Chart label** area, you can specify a textual label for the marker that is displayed within the chart.

Scale

These options are available for circular gauge and linear gauge charts.

Type and Orientation

> Change the **Type** of a circular gauge chart to determine whether the gauge be rendered as a quarter, half, or full circle, and which direction the gauge will be oriented in. Change the **Orientation** of a linear gauge chart to determine whether the gauge will be rendered horizontally or vertically.

Color Click **Custom Colors** to choose from a wider selection of colors or to create your own.

Center

These options are available for circular gauge charts.

Show Center
Click **Show Center** to display the center of the gauge. You can hide the center by deselecting this check box.

Color Select the color to use for the center. Click **Custom Color** to choose from a wider selection of colors or to create your own.

Needle

These options are available for circular gauge and linear gauge charts.

Style Change the **Style** of the needle to determine the shape of the needle on the rendered gauge. The available options are arrow, pointer, rectangle, tailed arrow, tailed pointer, and triangle. The default type is line.

Length
Change the **Length** of the needle to determine its span across the rendered gauge. This is displayed as a sliding scale.

Width Change the **Width** of the needle to determine its width on the rendered gauge. This is displayed as a sliding scale.

Color Select the color to use for the needle. Click **Custom Color** to choose from a wider selection of colors or to create your own.

Tick

These options are available for circular gauge and linear gauge charts.

Type Change the **Type** of the tick marks to determine their shape on the rendered gauge. The available options are circle, line, rectangle, reverse triangle, and triangle. The default type is line.

Tick Mark Color
Select the color to use for the tick marks. Click **Custom Colors** to choose from a wider selection of colors or to create your own.

Label You can click ☑ **Show Tick Labels** to turn off the range number that is displayed next to the tick mark. You can also select the color to use for the labels. Click **Custom Color** to choose from a wider selection of colors or to create your own.

Range

These options are available for circular gauge and linear gauge charts.

Attribute
Select the attribute to which the specified range applies. If you chose multiple attributes for the view, you can select another attribute from the ▾ list.

Range Values (Minimum Value and Maximum Value)
Minimum Value shows the default lowest possible value for the selected attribute. You can edit the values to reflect a more realistic range, if

needed. Select **Use Default** to return to the original values. **Maximum Value** shows the default highest possible value for the selected attribute. Specify a more realistic range, to decrease the data points in the gauge. Select **Use Default** to return to the original values.

Edit Subranges
> You can add subranges for the range specified by clicking **Edit Subranges...** and defining a subrange by following the same procedure used to define the range.

Apply to all views that use this attribute
> ☐ **Apply to all views that use this attribute** lets you apply this range, if appropriate, to this attribute on all other managed systems. The range is used only for gauge views of this attribute.

Add Overlay

After clicking the Plot Area of the chart's thumbnail graphic, the Bar Chart, Plot Chart, or Area Chart tab is displayed. This tab includes the **Add Overlay** button for creating a line chart overlay on the current chart.

Use this button to apply a line chart that overlays the plot area of the chart.

Attributes

The Attributes tab has the required fields for assigning the attributes to plot in the overlaying line chart, one attribute per line.

Attributes
> Select an attribute from the Available list and click ◀ to move it to the assigned list. You can use Ctrl+click to select multiple attributes or Shift+click to select all attributes from the first one selected to this one.
>
> Each attribute will have a different line color and can be styled to change the line type, line color, and legend symbol. To change the styling of the line chart, click an attribute in the **Assigned** list. This enables the **Attribute Style** fields, described next. After modifying the style for an attribute, repeat for any other assigned attributes.

Text The attribute name will be displayed in the legend label, or you can write something else here.

Line Style
> The line style is a solid 3-pixel wide line, or you can select one of the broken line styles from the **Pattern** ▾ and another thickness from 1 to 10 from the **Width** ▾ list. The overlay line width is 3 pixels or you can select any value from 1 pixel to 10 pixels. Use the **Color** ▾ list to see and select a different color for the overlay and its associated legend item, or click **Custom Colors** to select or create a new color to represent the attribute.

Symbol Style
> The default shape that will appear in the legend is a disc, or you can select another from the **Shape** ▾ list. The size of the shape is 5 pixels or you can select another size from the **Size** ▾ list. Use the **Color** ▾ list to see and select a different color for the shape, or click **Custom Colors** to select or create a new color.

Value Axis

Value Axis

○ **Use existing Value Axis** disables all other fields in this tab and the Axis Label tab that would otherwise be applied to the new value axis.

◉ **Use new Value Axis** creates a new value axis for the overlay line chart. The Axis Parameters settings are applied to the new value axis, as are the settings in the Axis Label tab.

To see where the new value axis would be displayed in the chart, position the Add Overlay window so that you can see the **Preview** area of the Properties editor and then select the other radio button.

Axis Parameters

☐ **Show Grid** is used to apply grid lines to the plot area background for the overlay values.

Font is the font to use for the axis values; **Size** is the font size in points; and **Style** is the font style applied to the axis values: Plain, Bold, Italic, or Bold & Italic.

◉ **Horizontal** applies the overlay axis values horizontally.

○ **Vertical** rotates the values to display them vertically along the overlay axis.

Axis Label

Label ☐ **Show Text** is selected to apply a label to the overlay value axis.

Text is the label for the overlay value axis.

Font is the font to use for the label; **Size** is the font size in points; and **Style** is the font style applied to the label: Plain, Bold, Italic, or Bold & Italic.

Orientation

◉ **Horizontal** applies the axis label horizontally.

○ **Vertical** rotates the label to display it vertically along the overlay axis.

Legend Label

Select the Legend Label area of the thumbnail graphic in the Properties editor
𝐀ₐ **Style** tab to control the display of the chart legend and its appearance, including the coloring of the data series: pie slices, bars, plot lines, and plot areas.

Show Legend - General

Use this tab to specify whether or not to display a legend for the chart, its placement, and its appearance. This tab is not available for gauge charts.

☑ **Show Legend**

Displays the legend with the chart.

☑ **Show Border**

Draws a border line around the legend box.

Position

The placement of the chart legend in the view. The default position is the **Right** side of the view. You can select a different position from the ▾ list.

□ **Place Legend in collapsible panel**
   When selected, creates a separate bordered pane for the legend inside the
   view that you can scroll, collapse, or drag to adjust the width.

□ **Trim Legend Labels**
   When selected, trims the legend labels to the boundaries of the legend box
   or panel in the manner specified in the **Trim Style** ▾ list.

Text   The **Text** area determines the font to use, its size and whether any styling
   is applied. Select any changes from the ▾ list for **Font**, **Size**, or **Style**.
   When ☑ **Use CSS values** is selected, this area is disabled.

☑ **Use CSS values**
   This is set to use the values of the workspace presentation cascading style
   sheet (`ws_pres.css`) for the label text (header, footer, and chart legend).
   This option is useful for text in double-byte character set (DBCS) languages
   that are more legible with a different font styling than the defaults. The
   Tivoli Enterprise Portal administrator can edit the `ws_pres.css` style sheet
   in the *Install_dir*`\cnb (Windows); cw (Linux)\classes\candle\fw\`
   `resources\styles` directory on the computer where the Tivoli Enterprise
   Portal Server is installed, to establish a specific font style that will be used
   for all label text and that will not be overwritten whenever your
   installation is upgraded to a new release.

   Clear the check box if you want to use the font styling that is specified in
   the **Text style** area.

Legend Label

Use this tab to specify the colors and text associated with the data points and data
series in the chart.

Legends
   Lists the attributes that are being plotted. Select an attribute from the list to
   see and change the legend text or the plot point (or series) color.

Text   The legend label text. By default, this is the attribute name.

Fill Style
   For pie charts, bar charts, and area charts, you can select a **Fill Pattern** or
   Color from the list for the data point and its associated legend item. If you
   want more color choices, click **Custom Colors** and select from a broader
   range of colors or create a new color.

Line Style
   For plot charts, you can select a line pattern and width for each attribute's
   line on the chart. You can also select the color to use for the lines. Click
   Custom Colors to choose from a wider selection of colors or to create your
   own.

Symbol Style
   For plot charts, you can select a symbol and size for each data point on an
   attribute's line. You can also select the color to use for the symbols. Click
   Custom Colors to choose from a wider selection of colors or to create your
   own.

Custom Color
The Custom Color window enables you to customize the color to exactly what you
want.

Swatches

Select a color in the Swatches tab or create a new color in the HSB or RGB tab.

HSB HSB displays a color box and vertical slider bar for adjusting the Hue and Saturation or Brightness and RGB values. Click the H, S or B radio button, and drag the small selection ring through the color box:

- As you move from left to right, Hue goes from 0 to 360 degrees and Saturation goes from 100% down to 0%
- As you move from top to bottom, Brightness diminishes.

Notice the values of the two *unselected* elements change as you drag the selection ring, but not that of the one selected. Drag the vertical slider bar up and down to change the value of the selected element.

RGB The RGB tab has horizontal sliders to adjust the Red, Green and Blue values.

Value and Category Axes

Edit the axis style properties to switch between vertical and horizontal orientation, and to add labels and styling options.

The value axis is arranged vertically and the category axis is arranged horizontally with tick marks and numbers at regular intervals to represent the full range of attribute values being plotted. Click the Value Axis or Category Axis of the bar chart, plot chart, or area chart thumbnail graphic to see the General, Axis Label and Value Axis tabs:

General

Value Axis - General tab and Category Axis - General tab:

☑ **Show Grid**

The axis already draws a line for the unit measurement, and you can cross these lines with grid lines to aid the reader in following the axis tick mark values.

☐ **Invert Category Axis and Value Axis**

Switches the position of the value axis with the category axis so the values are written horizontally along the bottom of the chart and the category (refresh interval times) axis is vertical.

Label

Axis Label tab:

Label ☑ **Show Text** displays a label for the category axis or value axis (such as "Bytes per second""). Type the label in the **Text** field. Select any formatting changes from the ▾ list for **Font**, **Size**, or **Style**.

Orientation

The label can be written horizontally or vertically. Choose the orientation that fits best in the space.

Axis

Value Axis tab and Category Axis tab:

Text The **Text** area determines the font to use, its size and whether any styling is applied. Select any changes from the ▾ list for **Font**, **Size**, or **Style**.

Orientation

> The text for the axis tick marks can be written horizontally, diagonally, or vertically. You can select each radio button and see in the chart **Preview** area how the labels are displayed with that orientation. You also have the slider bar for adjusting the rotation angle from 0° to 180°.

Footer Label

Select the Footer Label area of the thumbnail graphic in the Properties editor 𝐀ₐ **Style** tab to add a label at the bottom of the view.

☐ **Show**

> Select the check box to display the contents of the **Text** field at the bottom of the view.

Text   This is the text to display as the footer if you selected ☑ **Show**.

Text style

> Select a different **Font name** or font **Size** from the ▾ list. ☑ **Bold** makes the footer text boldface. ☐ **Italic** makes the footer text italicized. If ☑ **Use CSS values** is selected, this area is disabled.

☑ **Use CSS values**

> This is set to use the values of the workspace presentation cascading style sheet (ws_pres.css) for the label text (header, footer, and chart legend). This option is useful for text in double-byte character set (DBCS) languages that are more legible with a different font styling than the defaults. The Tivoli Enterprise Portal administrator can edit the ws_pres.css style sheet in the *Install_dir*\cnb (Windows); cw (Linux)\classes\candle\fw\ resources\styles directory on the computer where the Tivoli Enterprise Portal Server is installed, to establish a specific font style that will be used for all label text and that will not be overwritten whenever your installation is upgraded to a new release.
>
> Clear the check box if you want to use the font styling that is specified in the **Text style** area.

Chart baselines

Use the values from the current bar chart, plot chart, or area chart to perform predictive and trend analysis of threshold conditions. You can then anticipate problems or predict future values by comparing current values with past results, or overlay the chart with situation comparison criteria.

Related tasks:

"Modeling a situation with view data" on page 290
Use the values from the current chart or table to model threshold conditions and then use the results in a new situation. Historical views are particularly valuable for researching thresholds.

Adding monitored baselines to a chart

Apply a monitored baseline of a situation's threshold criteria to a bar chart, plot chart, or area chart to see how effective the situation is in this environment. You can visually compare the current values in the chart with the situation and adjust the thresholds as needed.

Before you begin

Monitored baselining works by comparing the values in the chart with the comparison criteria of a situation. You must have a situation defined with at least

one of the attributes in the chart and the situation distributed to the managed system (or managed systems) assigned to the workspace.

The situation is ineligible for monitored baselines if the formula has embedded or correlated situations or uses a group function (such as count or sum), in which case the Add Monitored Baseline window does not open when you click **Add Monitored Baseline**.

The monitored baseline tool does not show in the view toolbar if the numeric values in the chart are the result of a calculation. An example is the *Open Situation Counts - Last 24 Hours* chart in the Enterprise Status workspace.

About this task

Take these steps to add a monitored baseline to the chart.

Procedure

1. Click **Add Monitored Baseline** in a bar chart, plot chart, or area chart. The Add Monitored Baseline window opens with a list of eligible situations to select from.
2. Click the situation that you want to apply baselines from. Use Ctrl + click to select multiple situations, Ctrl + click a selected situation to deselect it, and Shift + click to select all situations between the one previously selected and this one.
3. If the situation has expression overrides (fx icon in the **Overrides Exist** column) and you want to see which threshold (the original or the override) is used at the sampling time, select the □ **Show effective thresholds only** check box. See "Monitored baselines chart examples" on page 128.
4. To apply the selected baseline or baselines to the chart, click **Apply** to keep the dialog open or **OK** to close it.

Results

The comparison criteria of the situation is displayed as a line or a pattern behind the plotted values. Any situation override thresholds are indicated by broken lines during the scheduled time when they apply.

What to do next

You can perform any of these activities in the monitored baseline chart:
- Move the mouse pointer over a baseline to highlight it and see the situation formula.
- Move the mouse pointer over a plot point to see the value, the name of the attribute, and when the data sample was taken.
- Click and drag to select a smaller time range to zoom in on, and then press Esc to zoom out.
- Right-click a monitored baseline and click **Edit Situation** to launch the Situation editor; or click **Edit Override** to launch the Expression Override editor.
- Right-click a plot point and click **Model Situation** to open the Model Situation chart and populate the attributes table with the values from the plot point. After performing additional analysis, you can create a new situation with the calculated values and edit them as needed.

- Remove the monitored baseline from the chart: Click ⊞ **Add Monitored Baseline** and Ctrl + click a situation.
- Apply another type of baseline if this is an area chart or plot chart: ⌁ **Add Statistical Baseline** and, for plot charts only, ⊞ **Add Historical Baseline**.

Watch the Add Monitored Baseline video to learn how you can add situation thresholds in a chart.

Related reference:

"Bar chart" on page 109
Add a bar chart to a workspace to compare values among related attributes. The stacking bar chart best shows multiple values for the same attribute. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Plot chart" on page 110
Add a plot chart to a workspace to show trends over time among related attributes. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Area chart" on page 112
Use the area chart for showing trends over time among related attributes. The area chart is like the plot chart except that the area below the plotted line is filled in with color to indicate volume.

Monitored baselines chart examples:

These are examples of charts with monitored baselines applied.

Plot chart with baselines applied from two situations
> This chart shows ✳ Disk Writes Per Second over the past 12 hours. The ▬ red baseline is from a situation with a threshold of 150 disk writes per second; the ▬ blue baseline is from a situation with a threshold of 60 disk writes per second. If these situations were started for this managed system, the "red" situation would have been true twice in the past 12 hours and the "blue" situation would have been true 5 times.
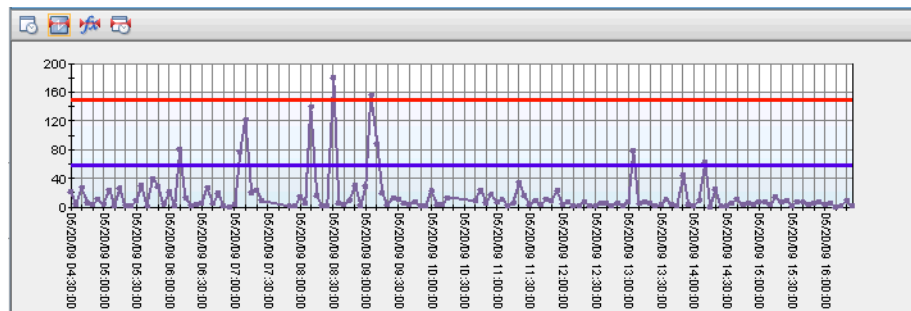
Plot chart with a baseline from one situation and its expression override
> This charts shows the Interrupts Per Second. The monitored baseline is from a ▬▬▬ situation with a threshold of 1000 interrupts per second or higher. The situation has an ▬ ▬ ▬ expression override of 1200 interrupts per second that is scheduled to run from 07:30 to 16:30 daily.
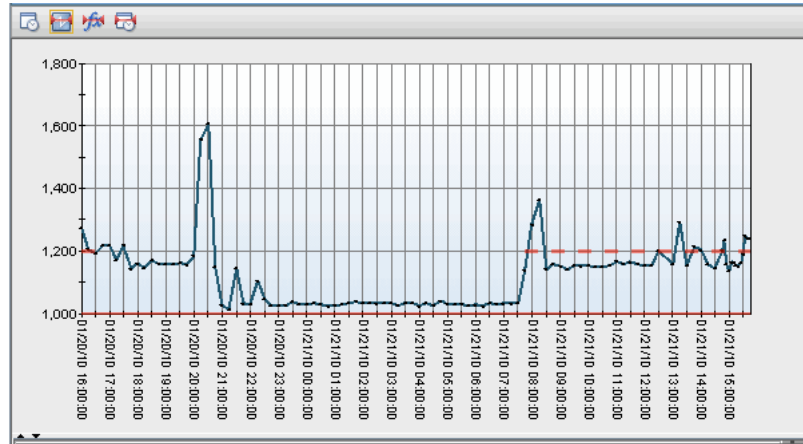
This is the same chart after ☑ **Show effective thresholds only** was selected. During the times when the expression override is applied, only that baseline is drawn. Connector lines are ——— solid to indicate that the threshold changes to the situation expression or – – – broken to indicate that the threshold changes to the override expression.
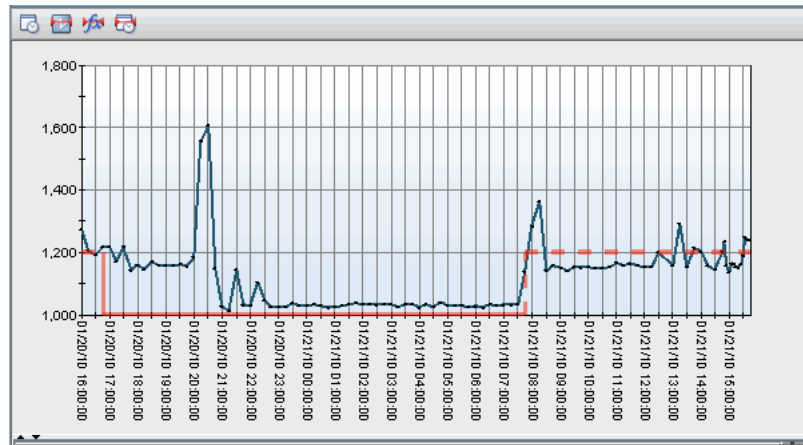
The baselines tell us that events are almost constantly opened for the situation except when the override is applied. The spike of 1600 interrupts per second is not noticed among all the other events.

The situation and the override were edited so that events occur only when the interrupts per second are higher than 1500 during off-peak hours and higher than 1250 during peak hours (between 08:00 and 20:00):
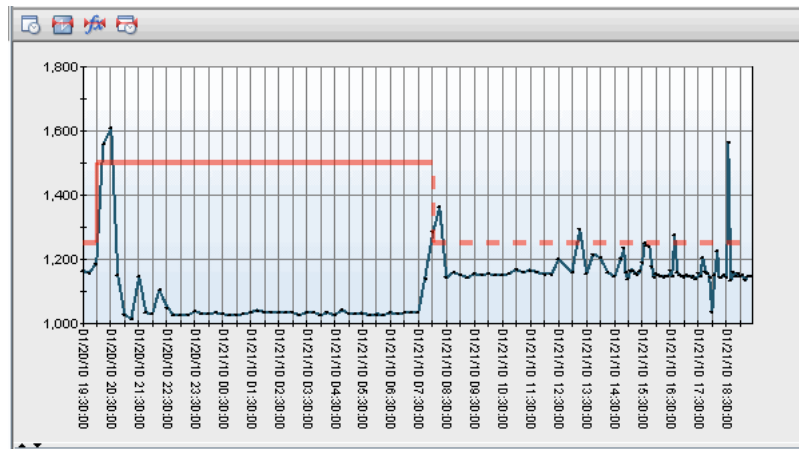
Rather than creating a custom schedule for peak hours, it might be more convenient to use the **PrimeShift** schedule, which runs from 08:00 to 17:00.

Related tasks:

"Adding monitored baselines to a chart" on page 126
Apply a monitored baseline of a situation's threshold criteria to a bar chart, plot chart, or area chart to see how effective the situation is in this environment. You can visually compare the current values in the chart with the situation and adjust the thresholds as needed.

Related reference:

"Bar chart" on page 109
Add a bar chart to a workspace to compare values among related attributes. The stacking bar chart best shows multiple values for the same attribute. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Plot chart" on page 110
Add a plot chart to a workspace to show trends over time among related attributes. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Area chart" on page 112
Use the area chart for showing trends over time among related attributes. The area chart is like the plot chart except that the area below the plotted line is filled in with color to indicate volume.

Adding statistical baselines to a chart

Use the statistical baseline function on a plot chart or area chart to apply any of these statistical functions to the selected attribute: Range between minimum and maximum, average, minimum, maximum, percentile, and mode.

Before you begin

The baseline function does not appear in the chart view toolbar if the chart values are derived (result of a calculation, such as the Enterprise Status workspace view, Open Situation Counts - Last 24 hours), if this is a plot chart with ⊙ **Attribute(s) across multiple rows** selected, or if historical data collection has not configured for the attribute group and started on the managed system assigned to the Navigator item. (You should see ▣ **Time Span** in the chart toolbar.)

For historical data beyond the past 24 hours, you must have a Tivoli Data Warehouse configured and the warehouse proxy agent that transfers data to it.

About this task

Take these steps to add a statistical baseline to the chart.

Procedure

1. While displaying a plot chart or area chart, select ⚑ **Add Statistical Baseline** in the view toolbar.
2. Select a check box to enable the function.

Function	Action
☐ RANGE - MIN/MAX	For the selected attribute, calculate the minimum and maximum values within the time period, and then shade that area on the chart.
☐ AVG	For the selected attribute, calculate the average value from the samples taken during the time period, and add or subtract the standard deviation that was entered in the **Arguments** field (valid values are -6 to 6).
☐ MIN	For the selected attribute, find the lowest sampled value within the time period, and increase or decrease that value by the percentage that was entered in the **Arguments** field (valid values are -100 to 100).
☐ MAX	For the selected attribute, find the highest sampled value within the time period and increase or decrease that value by the percentage that was entered in the **Arguments** field (valid values are -100 to 100).
☐ PERCENTILE	For the selected attribute, find the value below which NUM percent of the data points fall; where NUM is the value from the **Arguments** field (valid values are 1 to 100).
☐ MODE	For the selected attribute, find the most frequently observed value within the time period. No argument is entered for a MODE calculation.

3. Enter an **Argument**, if applicable, and click **Apply** to see the baseline in the **Result** cell and in the chart.
4. Try out different calculates by selecting other functions, selecting another attribute from the ▾ list, or selecting a different time span.
5. Click **OK** to close the dialog box, and perform any of these actions:
   - Move the mouse pointer over a baseline to highlight it; move the mouse pointer over a plot point to see the value, name, and time stamp.
   - Click and drag a smaller time range to zoom in, then press Esc to zoom out, once for each time you zoomed.
   - Right-click a data value that falls outside normal statistical baseline and click 🖳 **Model Situation** to open the Model Situation chart and populate the attributes table with the values from the data point. After performing additional analysis with the modeling tool, you can create a new situation with the calculated values and edit them as needed.

Example

Assume you have a chart of ✳ Disk Writes Per Second, but you want to see range over the past 2 hours. In the Add Statistical Baseline window, you select the **Last** 2 hours for the **Time Span**, and then select the ☑ PERCENTILE function with an argument of 95 to get a baseline of 95% over the past 2 hours. After you click **OK**, a − baseline is drawn, below which 95% of the data points fall. You can easily see the frequency of spikes (7 in the example) within the time frame.
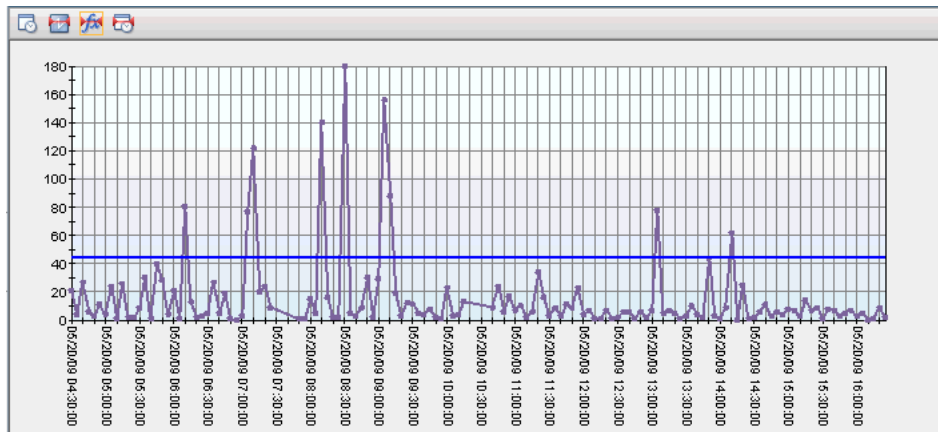
What to do next

After you have adjusted the baseline calculations and the time span to best illustrate the metric you are watching, you can save the workspace, and then open it whenever you want to observe the same baselines with current and recent data samples.

Watch the Add Statistical Baseline video to learn how you can determine when a metric is outside of its normal range for a system.

Related tasks:

"Modeling a situation with view data" on page 290
Use the values from the current chart or table to model threshold conditions and then use the results in a new situation. Historical views are particularly valuable for researching thresholds.

Related reference:

"Plot chart" on page 110
Add a plot chart to a workspace to show trends over time among related attributes. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

"Area chart" on page 112
Use the area chart for showing trends over time among related attributes. The area chart is like the plot chart except that the area below the plotted line is filled in with color to indicate volume.

Adding historical baselines to a chart

Use the values from the current plot chart to identify daily trends in threshold conditions. You can then better anticipate problems by comparing current values with expected future values or with past values, and overlaying the chart with situation comparison criteria, one plotted line for each day specified.

Before you begin

The ⊡ **Add Historical Baseline** tool does not appear in the plot chart tool bar if any of these conditions exists:

- Historical data collection has not been configured for the queried attribute group and started on the managed system (or managed systems) assigned to the Navigator item.
- A time span was specified. Click ⊡ **Specify Time Span for Query** and select the ◉ **Real time** to remove the time span.

- Any of the chart values are the result of a calculation.
- ⦿ **Attribute(s) across multiple rows** is selected. To see if it is, click ✎ **Edit Properties** in the view's toolbar, 🗛 **Style** tab, then the plot area of the thumbnail graphic.

For historical data beyond the past 24 hours, you must have a Tivoli Data Warehouse configured and the warehouse proxy agent that transfers data to it.

Procedure

1. Click ⊞ **Add Historical Baseline** in a plot chart. The tool does not display if historical data collection has not be configured for the attribute group (indicated by the presence of the ⊞ **Time Span** view tool).
2. To plot data samples taken yesterday, the same day last week, or a particular day or past days, select the check box for that time range.

   ☐ **Show yesterday's data**
   > Draws a plot line of data samples from the previous 24 hours. For example, if it is May 30 at 8:45 AM, the samplings will be from May 29, starting at 8:45 AM. The legend indicates the line color for each plotted attribute value from **Yesterday**. As new samples are added, the oldest are removed from the chart.

   ☐ **Show last week's data**
   > Draws a plot line of data samples from the same day last week. The legend shows the line color for each plotted attribute value from **Last Week**. For example, if this is Wednesday, the samples will be from last Wednesday.

   ☐ **Show this day's data**
   > Draws a plot line of data samples that have been taken on the day shown. If you would like to see a different day, click the date button to open a calendar. Select a month and day prior to today. For example, you had selected the check box to show yesterday's data, you could select the day before yesterday in the calendar to show the past two days.

   ☐ **Show data for last __**
   > Draws a plot line for each of the last number of **weekdays** or **calendar day**s you specified. For example, if today is Saturday and you select the last 7 calendar days, today and all days through last Sunday are plotted. The last 7 weekdays would be Monday through Friday of this past week plus the Thursday and Friday of the week before.

   Time Column
   > The time stamp type displayed will be shown along the X-axis. If the attribute group has other time stamp attributes, you can select one from the list. Not all time stamps will be displayed along the axis because of space limitations.
3. Click **Apply** to render the chart with a baseline for each chosen time range. The time ranges chosen are listed in the legend along with the color of the plot line, one line for each day.
4. Click **OK** to close the dialog box, and perform any of these actions:
   - Move the mouse pointer over a baseline to highlight it; move the mouse pointer over a plot point to see the value, name, and time stamp.
   - Click and drag a smaller time range to zoom in, then press Esc to zoom out, once for each time you zoomed.

- Right-click a data value that falls outside normal statistical baseline and click ⊞ **Model Situation** to open the Model Situation chart and populate the attributes table with the values from the data point. After performing additional analysis with the modeling tool, you can create a new situation with the calculated values and edit them as needed.

Example

The chart plots Disk Writes Per Second over a 12-hour period. We want to compare today's data, ☀ Monday, May 22, with the same period from the last two weekdays, ● Thursday, May 18, and ● Friday, May 19. In the Add Historical Baseline dialog box, select **Show this day's data** and **Show data for last 2 weekdays**. After you click **OK**, the two additional days are plotted. You can move the mouse pointer over a plot point to see the value and the date and time the sample was taken.
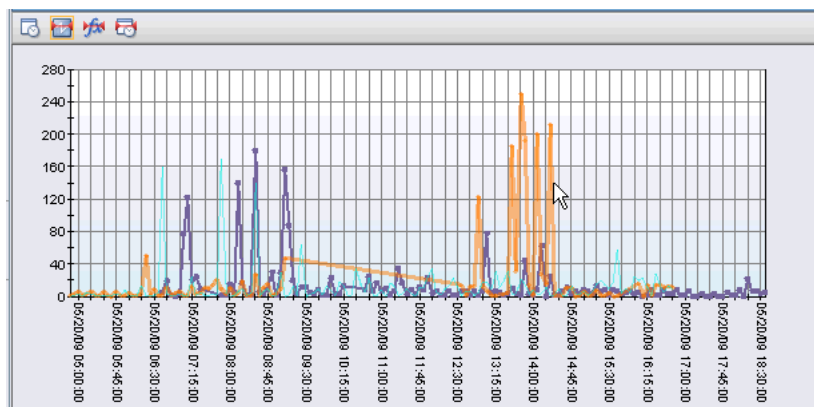
What to do next

After you have adjusted the baseline calculations and the time span to best illustrate the metric you are watching, you can save the workspace, and then open it whenever you want to observe the same baselines with current and recent data samples.

Watch the Add Historical Baseline video to learn how historical baselines can be used in line charts for either trend or predictive analysis.

Related tasks:

"Modeling a situation with view data" on page 290
Use the values from the current chart or table to model threshold conditions and then use the results in a new situation. Historical views are particularly valuable for researching thresholds.

Related reference:

"Plot chart" on page 110
Add a plot chart to a workspace to show trends over time among related attributes. Use the chart styling options that are available in the Properties editor to refine the appearance and to add options such as chart overlays.

Topology view

The topology view shows the arrangement of monitored components associated with its Navigator item. For attributes that can show their relationship to other attributes in the same group as connected objects, there is also the *relational table based* topology view.

As with the table view and chart view, you can use the topology view to display data from a query of your monitoring agents or from a custom ODBC query of a data source. The topology view provides a set of styling tools that enable you stress hierarchical relationships among related objects.

When you create a topology view that is based on a query (relational table based), each row of data returned becomes an object in the topology view. As with other query-based views, you can use the Properties editor to apply filters and thresholds to the view but with some differences:

- In the Filters tab, the column check box determines whether the value is included in the hover help. Any formulas you write for the filters are applied in the same way as for other query-based views.
- In the Thresholds tab, the chosen severity is applied to the topology as an alert indicator when the threshold is met.

Creating a topology view
About this task

A topology view can be added to a workspace to show the arrangement of monitored components associated with its Navigator item. Do the following to create a topology view.

Procedure

1. Select the workspace where you want to create the view. The context of the Navigator item and the queries available determine the type of information that can be displayed. For example, when adding a topology view to an Enterprise Navigator item workspace, you can choose Tivoli Management Services Infrastructure as your topology source or Relational Table Based. At no other level of the Navigator do you get the TMS Infrastructure source.

2. If you want the view to occupy a new space, click ⊞ **Split Vertically** or ⊟ **Split Horizontally** in an existing view. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty. If you do not see the ⊞ or ⊟ tools, the Split Pane Toolbar has been disabled. Select it from the View menu to enable it. You can also right-click inside a view and select ⊞ **Split Vertically** or ⊟ **Split Horizontally** from the pop-up menu.

3. Click ⛭ **Topology**, then click inside the space to adopt the new view.

4. Select a source type from the Select Topology Source window that opens. Relational Table Based is available at every level of the Navigator. Any product-specific topology sources will be unavailable from workspaces that are not associated with the product.

5. If the Properties editor opens, select 🔲 **Click here to assign a query** and follow these steps:

   a. Click ⊞ to expand a ⚙ monitoring agent and 🖥 attribute group branch.

   b. Select a 🔲 Query name to see its description and the specified attributes. The selected check boxes indicate which attributes are included in the query. The cells in the numbered rows show any filter criteria for the query.

   c. Click **OK** to assign the selected query to the view.

   The Preview window at the top of the Properties editor should show node icons and names. If not, select a column from the query for the topology object in the Column Mapping area of the Style tab.

6. If the Properties editor is not open, click ✎ Edit Properties in the view title bar. A view for which you have selected a query will be empty unless a column (attribute) from the query has been mapped to a topology object, which is done in the Style tab.

7. To limit the information in the object hover help, click the 🎋 **Filters** tab and clear any column check boxes whose values you do not want displayed. If you want to limit the number of objects in the view, you can also write a filter formula to include only the rows that satisfy the comparison.

8. To have a severity indicator applied to a topology object when a certain threshold has been exceeded, click the ▦ **Thresholds** tab, select an icon, and write a threshold expression for every severity to be represented.

9. To configure and format the view, click the A₀ **Style** tab and adjust the settings. The topology view has many styling options, so start with the essential ones when you first create the view:

   a. ▦ General to add the view title. This is also where you can limit the depth of the hierarchy if too much data is displayed.

   b. ▦ Mapping to specify the attribute to use for the nodes (Node ID and Name) and the attribute for the related nodes to point to or from (To or From Connection). Select the check box to create implied nodes and observe the change in the Preview pane to see if the view is improved by the selection.

   c. ▦ Layout to change the topology layout to the one that best suits the type of data relationships being presented.

10. When you are done configuring and formatting the view, save your changes with **Apply** to keep the Properties editor open or **OK** to close the editor.

What the topology view shows

A topology view added to a workspace maps the context of the Navigator item to a logical diagrammatic form. The objects shown are elements (attributes) of the product that is represented by the Navigator item. The elements of the monitored product are arranged to show their relationships to one another.

Manipulating the topology view

The view has tools for navigating the topology and arranging the objects to reveal areas of interest.

Toolbar

> In addition to the tools described here, the topology view toolbar for your product might have additional tools, such as an **Apply Pending Updates** button or **Status Filtering** button.

⬅ ⁻	⬅ Backward returns to the previous display of the topology. Click the ⁻ list button to see and select from a history. The view displays the selected content. Note that the tool does not show in the toolbar if your topology is small enough to fit within the view space.
➡ ⁻	If you went back to a previous topology display, click ➡ **Forward** to return to the most recently navigated version of the view. Click the ⁻ list button to see and select from a history. The view refreshes with the most recent status of the management elements.
▭	To zoom the display in and out, move the zoom slider switch right and left or press Num Pad + and -. As the zoom level decreases, the level of detail displayed is reduced. For example, the labels are automatically hidden.
⊡	⊡ Fit to view resizes the objects to fit into the current view space.
⊡	⊡ Overview opens a small window with the entire topology view in miniature. A blue frame indicates the visible area of the view space. As you scroll through the topology, the frame within the overview window adjust to show your location. Within the overview window, 🖐 click to move the viewpoint to that location, 🖐 point inside and ✛ drag the blue frame to move the viewpoint, and ↘ drag a handle of the blue frame to resize the view area.
▦ and 🖧	▦ View as Table switches the topology to table representation, one row for each object. Right-click the Node bar under Resource Type to open the pop-up menu. Note that if you are running the Thai language pack for Tivoli Monitoring, the topology table view is not translated but remains in English. Click 🖧 **View as Topology** to return to the topological display.
▭	Enter search text to locate the objects that match. As you type, the matches are automatically selected and zoomed.
➡ ⬅	When multiple matches are found, the **Find Next** and **Find Previous** buttons are enabled to allow you to navigate between the matches one object at a time. If your search text cannot be found, the background color will change to yellow. You can use * asterisk or ? question mark wildcard characters, such as **sys*** to find **system** and **sys_01**, or **system?** to find **system1** and **systemZ**.

Mouse and keyboard actions

> **Click**   Select the object. Use Ctrl + Click to select multiple objects. Click a corner of the area you want to select, drag to the opposite corner, then release the mouse button to select all the objects within or partially inside the boundary. Click outside the selection to cancel.
>
> **Drag**   Click an object, then drag to reposition it and the connectors to other objects.

Double-click

Perform the default action for your product on the object.

Right-click

Open the pop-up menu.

Hover Move the mouse pointer over an object to see its identifiers and status.

Alt + Click

(Not available on all UNIX platforms or configurations.) Move around the view with a virtual magnifying glass. Once the magnifier is displayed, you can release the Alt key.

Alt + L

Hide or display link labels.

Alt + N

Hide or display object labels.

Home Resize the objects to fit into the current view space.

Num Pad + and -.

Zoom in and out.

Popup menu options

Right-click a topology element to see the actions that can be applied to the selected object.

Right-click the topology view background to see the actions for the view as a whole.

Topology view style properties

The **Style** tab of the Properties editor has icons that you click to see and select the options for that aspect of the topology view: General, Node, View as Topology, Connection, Mapping, and Layout. The topology view for your product might have fewer options than these. For example, the Mapping and Layout options are not applicable to the TMS Infrastructure view, thus they are not included in the Style properties for that view.

General

Title The text entered here is displayed in the view's title bar.

Scalability

The topology view has an alternative table view display that can be selected from the view toolbar. You can also set a threshold number of objects that can be displayed in the topology before it switches automatically to a table.

Depth

Root Node enables you to identify a node to use as the root of the topology graph. When this field is empty, any node that does not have a connection TO it is treated as a root node. When a value is specified, the topology is displayed with the specified node at the beginning of the graph. Any objects that connect TO the specified node will not be displayed; only connections FROM the root node are visible. The value can be represented by a variable that is resolved through a workspace link. A variable is formed by enclosing the intended variable name in % percent symbols, such as %ROOT%.

Levels enables you to specify all levels of the topology graph or a specific number of levels, such as 1 to show only the root level.

Status Select the check box if you have created thresholds for the view and would like the node label to be highlighted to match the color of the severity indicator.

Mapping

This is where you identify the attribute to display as nodes and the attribute the nodes points to (or from).

Mapping

Node ID Column is the attribute that identifies the topology node.

Node Name Column is the attribute to use as the label for the topology object, usually the same attribute as the node ID.

To Connection Column, when chosen, becomes the object that the data nodes point to.

From Connection Column, when chosen, becomes the object that points to the data nodes.

☐ **Create implied nodes**, when selected, creates an object for an otherwise undefined connection in the topology. For some types of data, this item effectively unifies the topology.

Type **Node Type** is used to define the object type for a topology node:◉ **Default (kfw.TopoNode)** type for all topology nodes, ○ **Row type (<table>)** for the associated table row (value is is retrieved from the table definition and enclosed in parentheses), or ○ **Column** to define the object type, which allows you to select a different attribute for each topology object.

To Connection Type is used to define the object type for the connection to another topology node: ◉ **Default (kfw.TopoConnection)** for all *to* connections, or ○ **Column** to define the object type, which allows the attribute to be different for each *to* connection.

From Connection Type is used to define the object type for the connection from another topology node: ◉ **Default (kfw.TopoConnection)** for all *from* connections, or ○ **Column** to define the object type, which allows the attribute to be different for each *from* connection.

Connection

☐ **Show connection arrow**, when enabled, displays an arrowhead at the end of a connector.

Layout

The standard topology layout is `Tree`, or you can select one of several possible style structures.

- When the layout is set to **Tree**, you can adjust the flow direction, the positioning between objects (Mode), and how the connector lines are drawn.
- When the layout is **Hierarchical**, you can change the flow direction starting from the root object and the link style between objects at different hierarchical levels.

- When the layout is a **Grid**, use the options to set the maximum number of objects to display in each topology row or column.
- Select **Uniform Length Links** to have all connector lines the same length.
- Select **Spring Embedder** to vary the connector line lengths.

View as Topology

General

☐ **Disable direct object manipulations**, when selected, keep the objects in a fixed position that cannot be manipulated by the user. Normally you can move objects around the view, but the new position cannot be saved.

The **Change Color** button enables you to change the text (foreground) and highlight (background) color for a selected object.

Background

The **Change Color** button enables you to change the background color of the entire view.

Node

Text style fields enable you to adjust the label text style and color: Select another **Font name** from the list to change the font; select another point **Size** to increase or decrease the size of the font. Select the ☐ **Bold** check box to bold the label text; select the ☐ **Italic** check box to italicize the text. Click **Foreground Color** to change the color of the object label when it is selected.

Layout fields enable you to control the spacing and alignment of the label text line: **Wrapping Mode** is turned off by default. Select None for a single line of text; Word to wrap the text within the set line width; Truncate to truncate any text that extends beyond the set line width; or All to wrap text within the set line width and to truncate text that goes beyond the set line height. **Alignment** is Left, or can be changed to Center or Right. **Width** and **Height** define the label boundaries when text wrapping is set to Word or Truncate, as well as the point where Right alignment begins. The default value is 100 pixels.

Connection

Text style fields enable you to adjust the text style and color for the connection label. Click **Foreground Color** to change the color of the connector label when it is selected.

Layout fields enable you to control the spacing and alignment of the connection label line: **Wrapping Mode** is turned off by default. Select None for a single line of text; Word to wrap the text within the set line width; Truncate to truncate any text that extends beyond the set line width; or All to wrap text within the set line width and to truncate text that goes beyond the set line height. **Alignment** is Left, or can be changed to Center or Right. **Width** and **Height** define the label boundaries when text wrapping is set to Word or Truncate, as well as the point where Right alignment begins.

TMS Infrastructure view

The Enterprise Navigator item has a predefined workspace that includes a topology view of your Tivoli managed network called TMS Infrastructure. This view shows the relationships and linking of the Tivoli Management Services components.

Creating a TMS Infrastructure view
About this task

As well as the TMS Infrastructure view in the Managed System Status predefined workspace, you can add the same view to any Enterprise level workspace.

Procedure

1. Open the ▦ Enterprise level workspace where you want to create the view.
2. If you want the view to occupy a new space, click ▥ **Split Vertically** or ▤ **Split Horizontally** in an existing view. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty. If you do not see the ▥ or ▤ tools, the Split Pane Toolbar has been disabled. Select it from the View menu to enable it. You can also right-click inside a view and select ▥ **Split Vertically** or ▤ **Split Horizontally** from the pop-up menu.
3. Click ▰ **Topology**.
4. Click inside the space to adopt the new view.

Results

The hub monitoring server is shown with the components that connect to it arranged topologically. If, instead of a topology view, a table view is displayed, the number of objects to display exceeds the threshold.

Understanding the TMS Infrastructure view

The monitored network configuration and status information are collected and then arranged in the TMS Infrastructure topology view. In the initial base view, the Tivoli Management Services infrastructure of the hub monitoring server and portal server are displayed. If the infrastructure includes a warehouse proxy and summarization and pruning agent, or any remote monitoring servers, they will also be displayed. Double-click ⊞ Expand to drill down on a monitoring server and reveal all the agents reporting to it. Any offline components are indicated by a red alert indicator, such as ▧ for an offline monitoring agent. Move the mouse pointer over an object to see its status:

```
HUB_TIVOLIUSER
Online
Resource type:  TEMS
```

```
Product:  EM
Version:  07:10:04
IP Address:  9.55.105.185
Last heartbeat:  10/02/08 21:40:28
```

The status bar displays the total number of objects in the view and how many are
selected, and the date and time of the last view refresh.

The topology view can show all the monitoring components at once or show just
the *base topology* , which collapses the agent display under the local monitoring
server.

Monitoring agents subsumed under the remote monitoring server: ⊞

Monitoring agents subsumed under the hub monitoring server: ⊞

Event views

Besides relying on event indicators, you can use any of the event views to get
more granular information for analysis and action.

Event views present alternatives to event indicators in the Navigator. Message log
and situation event console views update automatically to show new events as
they arrive and changes in event status as they occur. Graphic views give you a
pictorial alternative to the Navigator for indicating alerts.

Adding a message log view

The message log gives you an overview of changes in situation status on your
monitored network. You can add the view to any workspace.

Before you begin

To create a new view and edit its properties, your user ID requires Workspace
Author Mode permission. For the message log view, your user ID also requires
View permission for Events.

Procedure

1. Open the workspace where you want to add the message log view.
2. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split
   Horizontally in one of the views. The pane divides and a duplicate is created. If
   you split the Navigator, the new view is empty.
3. Click ▤ **Message Log**. The mouse pointer changes to the view icon (hand icon
   in Linux). If you get the wrong icon or decide not to add the view, press
   Escape or click somewhere in the main toolbar.
4. Click inside the view where you want the message log.

Results

For each change of situation status in your monitored network, new rows are
added to the top of the view. Beyond the limit of 500 rows (apply quick settings to
lower the limit), old rows are deleted as new rows arrive.

What the view shows

The message log view displays a row of data for each situation status change. The status column includes indicators such as 🔴 **Stopped** and 🗙 **Deleted**.

As new rows arrive, they appear in descending order with the oldest rows at the bottom. After the maximum number of rows has been reached, the oldest are removed to make room for new rows.

Status   The status for the situation:

> ☑ **Acknowledged**: The event has been acknowledged.

> 🗙 **Deleted**: The situation has been deleted.

> ❗ **Problem**: The situation is in a problem state. Perhaps the managed system to which it is distributed is offline or is at an old, unsupported version level.

> 🛑 **Expired**: The acknowledgement of the event has expired while the event remains open.

> 🔲 **Opened**: The situation is running and is now true, which opens an event.

> ☑ **Closed**: The situation is running, was true and is now false.

> ↩ **Reopened**: The acknowledgement was removed before it had expired and the event is still open (*reopened*).

> ▷ **Started**: The situation is being started.

> 🔴 **Stopped**: The situation has been stopped.

Situation Name
> This is the situation display name and ID information. The message log view can be used to determine the display name associated with a Situation ID.

Display Item
> If the situation was defined with a *display item*, it appears here. Otherwise, this field is empty.

Origin Node
> The name of the agent and system where the situation is running.

Global Timestamp
> The time the event was recorded by the hub Tivoli Enterprise Monitoring Server.

Local Timestamp
> The time the event occurred or the condition was observed by the originating managed system. The format is MM/DD/YY HH:MM:SS, as in 07/23/11 14:45:35 for July 23,2011 at 2:45:35 PM.

Node   The name of the monitoring server from where the situation is monitored.

Type   The type of event. The possible values are:
> - **Sampled** for events that happen when a situation becomes true. Situations sample data at regular intervals. When the situation is true, it causes an event, which gets closed automatically when the situation goes back to false (or you can close the event manually).

- **Pure** for events that are unsolicited notifications. Examples of pure events are an out-of-paper condition on a printer and an occurrence that adds an entry to a log. Some monitoring agents have attributes that report pure events, such as the Windows Event Log and Windows File Change attribute groups. A situation using one of these attributes can monitor for pure events. Because of the nature of pure events, they are not closed automatically like sampled events; you must close the event manually. Alternatively, you can author the situation with an UNTIL clause.

Apply quick settings

Use the quick settings feature of the message log view to set a row limit, filter out unwanted rows, and highlight specific cell values.

Procedure

- Set a row limit
  1. Right-click anywhere in the message log view.
  2. In the popup menu, point to Set Max Rows, then select 50, 100 or 500 from the list. After the maximum is reached, the oldest row is removed. The view refreshes every time a new event arrives, so the more rows you allow in the view the more time the refresh takes.
- Set or remove a quick filter
  1. Right-click the cell on whose value you want to filter the incoming data.
  2. In the popup menu, select **Set Quick Filter**, then select the attribute (attribute name == value). If you want to remove all quick filters applied to this view, select Remove Quick Filters. You can add a quick filter to other cells to further refine the filtering. For example, you set quick filters to show only the ETE_eBA_Response_Critical situation when it has been raised.
- Set or remove a quick threshold
  1. Right-click the cell on whose value you want to apply a threshold indicator.
  2. In the popup menu, click **Set Quick Threshold**, then select the severity to apply to that value (attribute name == value); or if you want to remove all quick thresholds applied to the view, select Remove Quick Thresholds.

     Tip: If you edit the properties for this view, you can make the quick filter or quick threshold permanent by right-clicking the row number and selecting **Remove Quick Flag** in the **Filters** or **Thresholds** tab.

Editing the message log view properties

The properties for the message log view have **Filters**, **Thresholds** and **Style** tabs. You can set a filter so only certain status types or systems are reported, set thresholds to highlight values outside the limits, and define the appearance of this view.

About this task

With the message log view displaying, follow these steps to edit the properties.

Procedure

1. In the view title bar, click **Edit properties**. The Properties editor for the workspace opens and displays the properties of the selected view.
2. Set **Filters** to report only the values or value ranges you specify. For example, you might want to add a message log view to the Windows Systems

workspace that shows situation activity on the most critical systems. You could add filters to include only those systems. If you set a quick filter, it will appear here and the row will be marked with the ⚡ **Quick flag** icon. You can make it a permanent filter by right-clicking the row number and selecting ⚡ **Remove Quick Flag**.

3. Set Thresholds to highlight areas of interest. If you set a quick threshold, it will appear here and the row will be marked with the ⚡ **Quick flag** icon. You can make it a permanent threshold by right-clicking the row number and selecting ⚡ **Remove Quick Threshold**.

4. When you are finished editing the message log properties, save your changes by clicking **Apply** to keep the Properties editor open, or **OK** to close the editor.

Related tasks:

"Filtering a query-based view" on page 90
You can filter the information displayed on the view to show specific attributes and rows of data.

"Adding column thresholds" on page 101
You can decorate a cell in the table view (including event console, message log, and universal message views) with a background color, severity icon, or custom icon when the value exceeds a given threshold.

Adding a Universal Message Console view

Create a Universal Message Console view to see the messages produced at the hub Tivoli Enterprise Monitoring Server.

Situation and policy activities generate universal messages, such as when a situation has been created or deleted or a policy has been activated. The universal message console displays messages as they are received, one row per message.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want the view to appear.
2. If you want the view to occupy a new pane, click ▦ **Split Vertically** or ▤ **Split Horizontally** in one of the views.
3. Click ▧ **Universal Message Console**.
4. Click inside the view where you want the universal message console.

Results

After you click, the old view is replaced by the universal message console. The view is updated as messages are received by the hub monitoring server from the managed systems assigned to this Navigator item. Universal messages that are produced at other locations, such as a remote monitoring server or an agent, are not displayed.

Example

In this example, notice that the most recent entry shows a generated value rather than the situation name. In most cases the situation ID will reflect its display

name, but when the name contains non-English text or special characters such as *
asterisk and \ backslash, the situation ID will be a generated value.

Time Stamp	Message	Severity	Category
10/06/11 10:37:32	Enterprise situation Z00000008000000C8AA5E7D5119420F: HUB_MYSERVER is true	0	K041041
10/06/11 07:10:00	Enterprise situation Blue_alert_at_Rome:HUB_MYSERVER is true	0	K041041
10/06/11 3:50:56	User guest logged off server SRVR01	0	KDSMA009
10/06/11 3:50:45	Monitoring for enterprise situation NT_System_File_Warning started.	0	KO41046
10/06/11 3:50:44	Enterprise situation NT_Process_CPU_Critical: Primary:TIVOLIUSER:NT<Rvscan> is true	0	KO41041
10/06/11 3:46:11	Logon successful to server SRVR01 user guest:ip.pipe:#10.255.105.53(6015)	0	KDSMA011

What to do next

To match up a situation name with its ID, you can create a workspace with these
two views: Universal Message Console; Message Log. The message log view will
show the situation name with no restriction; the universal message console view
will show the situation ID as needed. You can also specify generation of universal
messages when creating a situation (Situation editor Action tab) or policy
(Workflows editor 🖼 **Take action or Write message** activity).

Adding a graphic view

The graphic view provides a canvas for you to cover with icons that represent the
Navigator items and that overlay the icons with alert indicators when events are
opened. You might create, for example, a graphic view showing your company's
floor plan or organization chart.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want to add the graphic view.
2. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split
   Horizontally in one of the views. The pane divides and a duplicate is created. If
   you split the Navigator, the new view is empty.
3. Click ▣ **Graphic View** The mouse pointer changes to the view icon (hand icon
   in Linux). If you get the wrong icon or decide not to add the view, press
   Escape or click somewhere in the main toolbar.
4. Click inside the view where you want the graphic.

Results

The old view is replaced by the graphic view with the default background and overlaid with icons representing the items within the current Navigator item branch.

What to do next

- Click ⍸ **Select**, then drag icons into position; use this and other tools in the graphic view toolbar to manipulate the view. See "Graphic view toolbar."
- If you want to add other Navigator items to the view, drag them from the Navigator view. You can drag items one at a time or use Ctrl+click to select multiple items, then drag them as a group. Be careful to drag and not click, which selects and opens the workspace for that Navigator item.
- If a new item appears on the Navigator after you created the graphic view, you can add it at any time by dragging its name from the Navigator view to the graphic view; if an item gets removed from the Navigator, you can remove its icon by selecting it with ⍸ ⌧ .
- If you want to change the background image or the style of the icons or their labels, edit the style properties. See "Graphic view style properties" on page 148.
- You can apply your own background image to the view. See "Creating a custom background image" on page 149.
- You can customize the view by changing the icons that display for the Navigator items to those of your choosing. See "Creating custom icons" on page 150.
- If you create multiple graphic views and change them to use your own background and Navigator item images, consider changing the default image set. See "Changing the default graphics" on page 152.

Graphic view toolbar

Use the graphic view toolbar to adjust the icons and the view's appearance.

Reload Stylesheets
    Resets the view and refreshes the icon style. It is useful if you are working with a view and the icons change state or move around unexpectedly.

Actual Size
    Displays graphics at 100%. Suitable for raster images such as .gif, .jpeg and .png files. Not applicable to map files and is disabled when a map background displays.

    If you clicked ⊡ Fit to Contents, Actual Size returns the background graphic and icons to their original size.

Fit to Contents
    Scales the view so that all icons are visible. This is useful when you have added or removed icons. Click Actual Size to go back to original size.

Zoom in
    Zooms in on the view to see an enlarged version.

Zoom out
    Zooms the view out to a smaller size.

Zoom Box
    Expands or contracts the view to fit a selected area.

Pan   Pans the image including background and icons.

Overview

Opens the overview window. The Overview window has a box, which you can grab to pan the image.

Select

Selection mode. This is a toggle tool: click it again to turn off selection. Use the Select tool to drag an icon into position or to select it for removal. You can click and drag from one point to the opposite corner to select a group of icons, or use Shift+click to select multiple icons.

Clear

Removes the icons that are highlighted with the Select tool.

Grid

Shows or hides the grid. This is a toggle tool: click it again to turn off the grid. It is especially useful as a guide for positioning icons on an even plane.

Undo

Voids the last action.

Redo

Repeats the last voided action.

Snapshot

Takes a snapshot photo of the viewable area. You can save the snapshot as a PNG (Portable Network Graphic) file by clicking ⬛ **Export** in the lower-right corner of the snapshot window. (See also the *Tip* under "Graphic view style properties.")

Print

Prints the active view.

Graphic view style properties

Use the Style tab or the Properties editor to change the appearance of the view background and the icons.

To edit the properties

1.

Click ◪ in the graphic view title bar.

2.

Click the plot area of the thumbnail graphic. On this tab you can select the background for the graphic view and specify a stylesheet for any text displayed.

Background

◉ **Map** to identify the background as a geo-referenced map.

○ **Image** to specify a background image of type PNG, JPG or GIF. Click **Browse** to open a list of images from which to choose or select <blank background> for no background image.

☐ **Fit to view**, when selected, automatically resizes the image to fit in the view space. If you adjust the borders of the workspace or graphic view, the image shrinks or expands to fit. This option is selected and disabled when **Map** is selected; cleared and enabled when **Image** is selected. If you leave it cleared, your image will retain its original sizing no matter what the size of the view.

Browse looks for maps on the Tivoli Enterprise Portal server computer under:

```
<itm_installdir>/cnb (cw for Linux)/classes/candle/fw/resources/
maps/users
```

Tip: If you are using a map for the background and do not want it resized as you adjust the borders of the view or workspace, create a snapshot image:

a.

Create the graphic view with the desired map and adjust the view and workspace borders until the map is the size you want and no white space shows.

b.

Click ⬚ **Select** and click and drag to select all the icons, then delete them with ✖ **Delete**.

c.

Click ⬚ **Snapshot**, then click ⬚ **Export** in the lower-right corner of the snapshot window to save the snapshot as a .PNG file.

d. Create a custom background and apply the image to the view.

Style The style sheet determines which icons to display for Navigator items and the format and positioning of their text labels. To change from the default, click **Browse** and select from a list of the available style sheets.

The icon used for the large and small icon style matches the icon of the Navigator item (such as ⬚ , ⬚ , ⬚ , ⬚ , and ⬚).

If the preview does not refresh as you edit, click **Test** to see how the view will look with your style changes. If the graphic view is showing behind the window, you can click **Apply** to see your changes in the full graphic view.

3.

When you are finished editing the graphic view properties, save them by clicking **Apply** if you want to keep the window open or **OK** to close the window.

Creating a custom background image

Create and apply your own non-georeferenced images for the graphic view background. Acceptable formats are PNG, JPG and GIF.

About this task

Complete these steps to apply an image of your choosing to the graphic view background.

Note: This release of Tivoli Enterprise Portal does not support the addition of geo-referenced maps.

Procedure

1. Create the image that you want to use as a background in a graphical editor and save as a .png, .jpg, or .gif file, using a one-word name (underscores are acceptable), and all lowercase if the Tivoli Enterprise Portal Server is on Linux or operating systems such as UNIX.

2. Copy the image to: *install_dir*\cnb (Windows); cw (Linux)\classes\candle\ fw\resources\backgrounds\user.

3. After adding the graphic view to a workspace, click ✐ **Properties**.

4. Click the middle (plot area) of the thumbnail graphic.

5. In the Background area, click ◉ **Image**.
6. If you want the image to shrink or expand to fit the full view space, select ☑ **Fit to view** If you leave this option cleared, the image will retain its original size; if it is larger than the view space, scroll bars are displayed; if it is smaller than the view space, it will be placed at the top left corner with white space fill-in.
7. Click **Browse**, and select the image that you created from the list of images under `background/users`.
8. Click **Apply** if you can see the graphic view is showing behind the window and want to see how your image will look in the view space.

Creating custom icons

Graphical icons are applied to the background image of the graphic view to represent the Navigator items. Customize the graphic view by applying icon graphics of your choosing rather than those from the Navigator view. For example, you might want to use your company logo instead of the ▦ provided icon to represent the Enterprise.

Before you begin

The procedure has you create the icon images, copy one of the Tivoli Enterprise Portal style sheets, edit it to reference your images, then update the graphic view style properties to use the new style sheet. The new style sheet and icons must be saved on the computer where the Tivoli Enterprise Portal Server is installed.

Create icons as close to 32 x 32 pixels (large_icon_) or 20 x 20 pixels (small_icon_) as possible. An image greater or smaller will be resized to those dimensions automatically and might look distorted. Also, the image will retain proper proportions only if it is the same height and width. For example, if your original icon is 76 x 54, resize it to fit 32 x 32 if you will be using one of the large icon style sheets.

Small icons are used in conjunction with the "small" stylesheets. The graphic itself should be no taller than 16 x 16 pixels with a 2-pixel top and bottom margin, for a total of 20 x 20 pixels. The 2-pixel margins ensure that the icon does not "touch" the one above or below it.

Procedure

1. Create the images:
   a. Create a one-word (all lowercase if the portal server is on Linux or UNIX) .gif or .jpg image file for every Navigator item icon you want to change in the graphic view. Use only ASCII characters in the filename. Names that use characters outside the ASCII code range will not be displayed in the list of available files when the user browses to select one. A set of large and small icons that you can use are located in the `icons\user` folder:

   b. On the computer where the portal server is installed, copy the images:
   Windows `install_dir\cnb\classes\candle\fw\resources\icons\user`
   Linux **UNIX** `install_dir/cw/classes/candle/fw/resources/icons/user`
2. Create a style sheet:

a. Change to the Navigator view whose graphic view icons you plan to change. Having the Navigator open while you type the item names in the style sheet ensures their correct spelling and letter casing.

b. Identify the Tivoli style sheet most like the one you want. It must be one of the large_icon_ or small_icon_ style sheets (not shape_ or parent_). The style sheet controls the color, size and position of the text labels (Navigator item names), as well as any icons you do not replace with custom icons.

c. Make a copy of the style sheet from the following directory, giving it a one-word name that identifies it as one that you customized, and place it in the child directory (user):

Windows *install_dir*\cnb\classes\candle\fw\resources\styles

Linux **UNIX** *install_dir*/cw/classes/candle/fw/resources/styles

d. Open your new style sheet in a text editor.

e. Edit the @*import* lines at the top of the file to insert ../ in front of the $. Example:

```
@import "../$base.css"; @import "../$baseLargeIcon.ss";
@import "../$baseNodeLabelBottom.css";
```

f. Copy and paste the following entry at the end of the file, one line for each Navigator item whose graphic view icon you want to change.

```
node[name = "My Item Name"]{iconImage: url(../../icons/user/iconname.gif);}
```

g. Replace `Your Item Name Here` with the exact name of the Navigator item; and replace `iconname.gif` with the exact name of the icon graphic file. Be careful to follow the syntax exactly. Any errors, such as missing or extra parentheses, quotes, or misspellings will give unpredictable results. The symptom might be that your style sheet change did not affect the graphic view or the view contents disappear.

h. Save the text file.

3. Apply the new style sheet:

a. Open the workspace that contains the graphic view or add a graphic view to the current workspace.

b. In the graphic view title bar, click ✐ **Properties**.

c. In the **Style** tab, click the plot area of the thumbnail graphic.

d. Adjust the **Background** settings as needed.

e. In the **Style** area, click **Browse**, select the custom style sheet from the **Select Name** list, and click **OK**. Your style sheets are displayed as user/style at the end of the list. If you do not see your style sheet in the list, check that you saved it in the right directory and that you did not start the name with $. If you are in browser mode, make sure you added a style entry to $_index.txt and that you spelled it correctly. You can drag the border between **Preview** and **Style** to make the preview image larger.

f. When you are finished editing the graphic view properties, save your changes with **Apply** to keep the window open, or **OK** to close the Properties editor.

What to do next

1. When you erase the background image name (no image selected) or use a non-map background, you usually need to save the workspace, switch to another workspace, then return to it to see the correct icon sizes. This is required whenever the image sizes do not scale as you expect.

2. After you have created a user style sheet and associated it with a graphic view, you can make further changes with a text editor and see the effect without

reopening the Properties editor. To do so, save the changed file, click ![icon] Reload Stylesheet to refresh the style sheet and see the new icons.

3. If you want your new style sheet to be applied automatically to all new graphic views you create, you can make it the default.

Changing the default graphics

When you drop a new graphic view onto the workspace, the background that displays is the default image, as are the icons overlaying the graphic. You can change the graphic and the styling through the Properties editor Style tab. If you prefer another background graphic or icons for the default images, you can use your own.

About this task

Take these steps on the computer where the Tivoli Enterprise Portal Server is installed.

Procedure

- To change the default background:
    1. Rename the original `default.jpg` to `default_backup.jpg` in the backgrounds directory:

       `Linux` or `UNIX` `install_dir`/cw/classes/candle/fw/resources/backgrounds

       `Windows` `install_dir`\cnb\classes\candle\fw\resources\backgrounds
    2. Copy the new default graphic to the same directory and name it `default.jpg`.
- To change the default icons and styles:
    1. Make a backup copy of `$base.css`, `$baseIcon.css`, `$baseLargeIcon.css`, and `$baseNodeState.css` in the styles directory:

       `Linux` or `UNIX` `install_dir`/cw/classes/candle/fw/resources/styles

       `Windows` `install_dir`\cnb\classes\candle\fw\resources\styles
    2. Edit the original style sheets to reference the images that you want as the defaults. All users connecting to this portal server will get this default icon images whenever they create a new graphic view.

Adding a situation event console view

You can add a situation event console view to list each alert for situations associated with the current Navigator item and those it subsumes.

Before you begin

This feature requires your user ID to have Workspace Author Mode permission.

About this task

The situation event console view adds a row for every event status change for situations associated with this branch of the Navigator, up to 100 rows. The console has a toolbar for filtering the view to show only the events that you want to see and a pop-up menu with items for managing events. Take these steps to add a situation event console view to a workspace:

Procedure

1. Open the workspace where you want to add the situation event console view.

2. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.

3. Click ▦ Situation Event Console. The mouse pointer changes to the view icon (hand icon in Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the main toolbar.

4. Click inside the view where you want the event console. After you click, the old view is replaced by the event console.

5. Use any of the filtering tools to limit or expand the rows to show the events you are most interested in.

Results

The view refreshes automatically, adding a row for every status change of the situations associated with this Navigator item and any items below it on the same branch. If you see no rows it means no events have be opened for any situations associated with this Navigator item.

What the situation event console shows

The Navigator alerts you when a situation becomes true by overlaying the Navigator item with an event indicator. Multiple alerts for the same Navigator item are consolidated under one event indicator, so you might not be aware of newly arrived alerts just by looking at the indicator. The situation event console view shows you the event activity for situations associated with the current Navigator item and any items within the branch.

You can navigate from an individual event to detailed information by clicking the link icon next to each event or by right-clicking the event and clicking **Situation Event Results**.

Toolbar

A row is inserted in the event console for every event action associated with the current Navigator item, one row per event. The event severity filters show a count of the events of that type next to the tool. For example, ⚠ 5 **Filter Minor** is a count of five minor events. Use the filter tools to display only events of a certain state or status.

⊗ **Filter Fatal**	Show only Fatal events.
⊗ **Filter Fatal**	Show only Critical events.
⚠ **Filter Minor**	Show only Minor events.
⚠ **Filter Warning**	Show only Warning events
✅ **Filter Harmless**	Show only Harmless events
ℹ **Filter Informational**	Show only Informational events.
◈ **Filter Unknown**	Show only Unknown events.
◧ **Filter Open**	Show Open events only, including any that were reopened (acknowledgement removed) or whose acknowledgment expired.
☑ **Filter Acknowledged**	Show Acknowledged events only.
🔴 **Filter Stopped**	Show Stopped situations only.
❕ **Filter Problem**	Show Problem situations only. These are situations that are in error for some reason. Use this filter to locate situations that are not functioning properly so you can diagnose the problem.
⊙ **Filter Expired**	Show the events that were acknowledged, the acknowledgement expired, and the condition is still true.

Find	Open the Find dialog box for finding events of a particular severity, status, or some other criteria. See "Finding table data" on page 48.
(Active) -or- (Paused)	The event console gets updated as status changes occur. Use this tool to pause automatic refresh, and click again to resume.
Total Events: *31* -or- **Filtered Events:** *7* of *31*	This is the current row count of events in the console and any events waiting to be added.
	If a filter has been applied, the count of events that are included in the filtered view is shown. For example, if you click the Fatal and Critical filters to show only the 2 Fatal events and 3 Critical events out of a total of 31 events, the display is Filtered Events: 5 of 31.
Item Filter: *Item_name*	The item filter is the name of the Navigator item by which the events are filtered.

Event list

The event console displays event information in the view.

	If you begin to scroll the view or expand pure events, automatic refresh turns off and this indicator is displayed at the top left corner of the list. Click **Console Resume** when you want to resume automatic refresh.
	Pure Event in the first column indicates the event is a pure event. Consecutive pure events are consolidated: Expand the list to show a row for each pure event; **Link** to event details about the pure event; Collapse the list back to one row. If the list does not expand, it means there is only one pure event.
	Tip: You can remove old pure events from the list by closing the pure event.
	The link button enables you to navigate to the Event Details workspace for an individual event.
Severity	The severity of the event: Critical, Warning, Informational, and so on.
	The status of the event:
	Open The situation is true and the event indicator turned on. The icon and threshold color indicate the severity of the alert.
	overlays the state icon if the acknowledgement for the event has *Expired* while the event is still open.
	overlays the state icon if the acknowledgement has been removed while the event is still open (*Reopened*).
	Acknowledged
Status	The event has been acknowledged.
	Stopped
	A user stopped the situation. It will not begin running again until someone starts it or edits it in the Situation editor.
	Deleted
	The situation was deleted through the Situation editor.
	Problem
	The situation is not functioning properly. One reason might be that the agent to which it is distributed is offline or is at an old, unsupported version level.

Owner	The user who acknowledged or took ownership of the event.
Situation Name	The name of the situation or policy.
Display Item	If the situation was defined with a display item, it is displayed here. Otherwise, this field is empty.
Source	The name of the agent and system identifying where the event occurred.
Impact	The name of the Navigator item where the event indicator originated; when you click ▼ , the entire branch opens.
Opened	This is the local time at the agent location when the event was last opened.
Age	This is the time interval that has passed since the event was opened. If you apply a Filter or Threshold to this column, enter the number of minutes. After you click outside the cell, you will see the value converted to days, hours and minutes. For example, an entry of `1450` becomes `1 Day, 10 Minutes`.
Local Timestamp	The time at the Tivoli Enterprise Monitoring Server (the one to which the agent connects) when the situation status changed. Additional information about the event, usually the event type.

Type

> **Sampled**
>
> > Sampled events happen when a situation becomes true. Situations sample data at regular intervals. When the situation is true, it causes an event, which gets closed automatically when the situation goes back to false (or you can close it manually).

> **Pure** Pure events are unsolicited notifications. Examples of pure events are an out-of-paper condition on a printer and an occurrence that adds an entry to the Windows System Log. Some Tivoli monitoring agents have attributes that report pure events, such as the NT Event Log and NT File Change attribute groups. A situation using one of these attributes can monitor for pure events. Because of the nature of pure events, they are not closed automatically like sampled events; you must close the event manually. Alternatively, you can author the situation with an UNTIL clause.

Situation ID	The situation identifier.

Table view features

The features available to the table view, such as column locking, can also be used in the situation event console view.

Pop-up menu

You can open the Situation Event Results workspace and perform other event related tasks from the pop-up menu: Edit Situation, Start Situation, Stop Situation, Acknowledge Event, and Event Notes. You can also launch an application on your system, link to another workspace, or send a command to the managed system where the event occurred.

Applying filters and thresholds

You can apply filters and thresholds to adjust the event list and the number of rows that are displayed from the console.

About this task

You can adjust the event list to show only those of a certain severity or status, owner, attribute or attribute value, or associated with another Navigator item. You can also limit or expand the maximum number of rows that can show on the console, and you can apply thresholds to cells. Tools are provided for the most common filters, state and status, so you can quickly toggle a filter on or off.

Severity

All severities are shown in the event list. You can filter the list to show one or two states instead of all:

- Click ⊗ , ⊗ , ⚠ , ⚠ , ☑ , ⓘ or ◈ to filter the view to show only events with that severity.

- Click a tool again to remove that particular filter.

Status All status types are shown in the event list. You can filter the list to show only a certain status:

- Click ▣ **Open**, ☑ **Acknowledged**, ⬤ **Stopped**, or ▮ **Problem** to filter on one or more of these statuses.

- Click a tool again to remove that particular filter.

Column

The toolbar enables you to filter the list by Severity, Status, or Owner. Through the popup menu, you can filter on any column and even multiple columns.

1. Right-click the cell on whose value you want to filter the incoming data.

2. Click ⚒ **Set Quick Filter**, then select the attribute (attribute name == value).

If you want to remove the quick filter, right-click a cell in the same column, point to ⚒ **Set Quick Filter**, and click either **Remove Quick Filter** to remove the last filter applied or **Remove Quick Filter Predicate** and select the filter to remove.

Navigator item

Initially, the event console shows events for the situations associated with the current Navigator item. You can change the view to show the event list for a different Navigator item:

1.

   Select the Navigator item whose events you want to see.

2.

   Hold down the Ctrl key while dragging the Navigator item into the view.

3.

   When the mouse is over one of the rows, release the mouse button.

Number of rows

Tivoli Enterprise Portal enters a row in the table at the time a change to the event status takes place. As new rows arrive, they appear in descending order with the oldest rows at the bottom. After the maximum (100 is the default) is reached, the oldest row is removed as each new row arrives. The view refreshes every time a new event arrives, so the more rows you have in the view, the more time the refresh takes. To change the maximum:

1. Right-click anywhere in the event console view.
2. Click **Set Max Rows**, then select 50, 100 or 500 from the list.

Quick threshold

Use the quick settings to highlight cell values that fall outside the threshold:

1.
   Right-click the cell on whose value you want to apply a threshold indicator.

2.
   Click   **Set Quick Threshold**, then select the severity to apply to that attribute value (attribute name == value).

If you want to remove the quick thresholds applied to the view, right-click a cell in the same column, point to   **Set Quick Threshold**, and click either **Remove Quick Thresholds** to remove the last threshold applied, or **Remove Quick Thresholds Predicate** and select the threshold to remove.

If you edit the properties for this view, you can make the quick filter or quick threshold permanent by right-clicking the row number and clicking Remove Quick Flag in the Filters or Thresholds tab.

Editing the situation event console view properties

The properties for the situation event console view have   **Filters**,   **Thresholds** and   **Style** tabs. You can set a filter so only certain status types or systems are reported, set thresholds to highlight values outside the limits in red or yellow, and define the appearance of this view.

About this task

With the situation event console view displaying, follow these steps to edit the properties:

Procedure

1. In the view title bar, click   **Edit properties**. The Properties editor for the workspace opens and displays the properties of the selected view.
2. Set   **Filters**, as described in "Filtering a query-based view" on page 90, to report only the values or value ranges you specify. If you set a quick filter, it will appear here and the row will be marked with the   **Quick flag** icon. You can make it a permanent filter by right-clicking the row number and selecting

     Remove Quick Flag.
3. Set   **Thresholds**, as described in "Adding column thresholds" on page 101, to highlight areas of interest. If you set a quick threshold, it will appear here and the row will be marked with the   **Quick flag** icon. You can make it a permanent threshold by right-clicking the row number and selecting

     Remove Quick Threshold.
4. When you are finished editing the situation event console properties, save your changes by clicking **Apply** to keep the Properties editor open, or **OK** to close the editor.

Adding a common event console view

Add the common event console view to a workspace to get a single, integrated display of events from multiple event repositories, including Tivoli Monitoring, Tivoli Enterprise Console, and Tivoli Netcool/OMNIbus.

About this task

To add a common event console view:

Procedure

1. Open the workspace where you want the common event console view. A common event console view is not associated with any particular Navigator item. Therefore, it always shows all available events from the configured connectors, subject to any filtering options you select for the view.
2. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
3. Click ▥ **Common Event Console**. The mouse pointer changes to the view icon (or the hand icon on Linux systems). If you accidentally click the wrong icon or decide not to add the view, press Esc or click somewhere in the toolbar.
4. Click inside the view where you want the common event console. After you click, the old view is replaced by the common event console. If any of the configured common event connectors are experiencing problems retrieving event data, a window might appear to notify you of the problem. You can use the Common Event Connector Status window to see the status of all configured connectors.
5. Use any of the filtering tools to show the events that you are most interested in.

What to do next

You can adjust the event console to show only events of a certain severity, status, owner, or other attribute value. You can also apply thresholds to cells.

What the common event console view shows

A row is inserted in the common event console for each event retrieved from a connected event repository: Tivoli Enterprise Monitoring Server, Tivoli Enterprise Console event server, and Netcool/OMNIbus ObjectServer. Use the tools to display only events of a certain severity or status, to find specific events, and to control how the display is updated.

Toolbar

When severity or status filters are applied, the event console shows the events that match any of the selected filters. If you filter on both severity and status, the console shows only the events that match both sets of filters. The exceptions are as follows:

- If you click all of the status filter icons, all status filters are removed; events of all status values appear in the console.
- If you click all of the severity filter icons, all severity filters are applied; events of all severity values appear in the console.

⊗	**Filter Fatal**	Show only Fatal events.
⊗	**Filter Critical**	Show only Critical events.

⚠ **Filter Minor**	Show only Minor events.	
⚠ **Filter Warning**	Show only Warning events	
☑ **Filter Harmless**	Show only Harmless events	
ℹ **Filter Informational**	Show only Informational events.	
◈ **Filter Unknown**	Show only Unknown events.	
◪ **Filter Open**	Show Open events only, including any that were reopened (acknowledgement removed) or whose acknowledgment expired.	
☑ **Filter Acknowledged**	Show Acknowledged events only.	
⬤ **Filter Stopped**	Show Stopped situations only.	
⚠ **Filter Problem**	Show Problem situations only. These are situations that are in error for some reason. Use this filter to locate situations that are not functioning properly so you can diagnose the problem.	
⬤ **Filter Expired**	Show the events that were acknowledged, the acknowledgement expired, and the condition is still true.	
🔍 **Find**	Open the Find dialog box for finding events of a particular severity, status, or some other criteria. See "Finding table data" on page 48.	
⏸ **(Active)** -or- ▶ **(Paused)**	The event console gets updated as status changes occur. Use this tool to pause automatic refresh, and click again to resume.	
Total Events: *31* -or- **Filtered Events:** *5* of *31*	This is the current row count of events in the console and any events waiting to be added. If a filter has been applied, the count of events that are included in the filtered view is shown. For example, if you click the ⊗ Fatal and ⊗ Critical filters to show only the 2 Fatal events and 3 Critical events out of a total of 31 events, the display is Filtered Events: 5 of 31.	
Item Filter: *Item_name*	The item filter is the name of the Navigator item by which the events are filtered.	

Event list

By default, the common event console shows the following information for each event:

Severity	The severity of the event. The possible values for a specific event vary depending on its originating event repository, and how the connectors map those severity values to those used by the console. Severity values include: • **Fatal** • **Critical** • **Minor** • **Warning** • **Harmless** • **Informational** • **Unknown**

Status	The status of the event. The possible status values for a specific event vary depending on its originating event repository. Status values include:
	• **Open**
	• **Reopened**
	• **Problem**
	• **Acknowledged**
	• **Closed**
	• **Deleted**
	• **Started**
	• **Stopped**
	• **Unknown**
	• **Response**
	If the originating event repository does not supply a status value, this cell in the table is blank.
	Note: Changes to event status might not be immediately reflected in the common event console. This is because the console retrieves some status data by periodically polling the originating repositories; changes to the status of an event are reflected the next time the console polls the repository. The polling interval is specified in the common event console configuration; the default interval is 1 minute.
Repository	The type of event repository from which the event originated. Currently, the supported values include the following:
	• **ITM** (IBM Tivoli Monitoring)
	• **TEC** (Tivoli Enterprise Console)
	• **OMNIbus** (Netcool/OMNIbus)
Name	The name of the event. The name is taken from one of the following values:
	• Situation name (Tivoli Monitoring events)
	• Event class name (Tivoli Enterprise Console events)
	• AlertGroup field (Netcool/OMNIbus events)
Display item	The display item for a Tivoli Monitoring situation event. This field is not used for events from other repositories.
Source	The host name or IP address of the monitored system from which the event originated.
Time created	The date and time when the event was generated.
Type	The Tivoli Monitoring event type (**Pure** or **Sampled**). For events from other repositories other than Tivoli Monitoring, this value will always be **Pure**.

In addition, the following fields are available but not shown by default (you can show them by modifying the properties of the view):

Connector name	The name of the common event connector that retrieved the event, as specified in the connector configuration.
Time Modified	The date and time when the event was most recently modified.
Owner	The owner of an acknowledged event, taken from one of the following values:

- **administrator** attribute (Tivoli Enterprise Console events)
- The user ID of the user who acknowledged the event (Tivoli Monitoring events)

If the event has not been acknowledged, this field is empty. This field is not used for Netcool/OMNIbus events.

Message	The message associated with the event. The message is taken from one of the following values:

- **msg** attribute (Tivoli Enterprise Console events)
- Summary field (Netcool/OMNIbus events)

This field is not used for Tivoli Monitoring events.

Incident ID	This field is not currently used.
Extra Column 1	These fields are provided for mapping additional event data. If these fields are being used to display user-defined columns, the field names might appear different from those shown here.
Extra Column 2	
Extra Column 3	
Extra Column 4	
Extra Column 5	
Short Name	

The features that are available to the table view, such as column locking, can be used in the common event console view. (See "Table view features" on page 100.)

Filtering events by column value

You can use the popup menu to filter the displayed events based on any displayed attribute value, or a combination of attribute values.

About this task

To apply a quick filter using the popup menu:

Procedure

1. Right-click in a table cell that contains the value you want to use to filter the displayed events. For example, if you want to see only events originating from a particular monitored system, right-click a cell in the **Source** column that contains the name of the system you are interested in.
2. Click **Set Quick Filter**.
3. Select the filter you want to apply. The initial filter is specified using the format *attribute* **EQ** *value*. For example, the filter **Owner EQ SYSADMIN** matches events whose owner is SYSADMIN.
4. To apply an additional filter, repeat the same steps using a different table cell. An additional filter is specified as a logical AND clause (for example, **AND Situation EQ NT_Log_Space_Low**).

Removing a quick filter from an event console

You can use the pop-up menu in any table cell to remove a quick filter, or to remove all quick filters.

About this task

To remove a quick filter:

Procedure

1. Right-click in any table cell.
2. Select **Set Quick Filter**.
3. Select one of the following:
   - To remove a single filter condition, select **Remove Quick Filter Predicate** and then select the filter predicate you want to remove.
   - To remove all quick filters, select **Remove Quick Filter**.

   Note: Removing filters using the pop-up menu does not affect any filtering applied using the toolbar icons.

Removing a quick threshold from an event console

You can use the pop-up menu in any table cell to remove a quick threshold or to remove all quick thresholds.

About this task

To remove a quick threshold:

Procedure

1. Right-click in any table cell.
2. Select **Set Quick Threshold**.
3. Select one of the following:
   - To remove a single threshold condition, select **Remove Quick Threshold Predicate** and then select the threshold predicate you want to remove.
   - To remove all quick thresholds, select **Remove All Quick Thresholds**.

Setting the view properties for an event console

The properties of an event console include filters, thresholds, and style settings. You can set a filter so only certain events or columns are displayed, set thresholds to highlight cells containing specified values, and configure the appearance of the view.

About this task

To configure the view properties:

Procedure

- Right-click anywhere in the event console view.
- Click ▦ **Properties**.
- Click the Filters tab to configure filters. Use this option if you want to configure the view to show only events that match certain values or value ranges.

If a quick filter has been set, it is displayed here, and the row in the table is marked with the ⚡ **Quick flag** icon. You can make the filter permanent by right-clicking the row number and selecting ⚡ **Remove Quick Flag**.

You can also use the Filters tab to select and order the displayed columns.

- Click the Thresholds tab to highlight areas of interest.

  If a quick threshold has been set, it is displayed here, and the row in the table is marked with the ⚡ **Quick flag** icon. You can make the threshold permanent by right-clicking the row number and selecting ⚡ **Remove Quick Flag**.

- Click the Style tab to add a header or footer to the view.

- When you have finished editing the event console properties, click **Apply** to apply your changes to the view, or **OK** to apply your changes and close the Properties window.

What to do next

Note: To keep the changes you have made, click 💾 **Save** to save the workspace.

Viewing common event connector status

You can use the toolbar to view the status of the configured common event connectors.

About this task

Common event connectors retrieve event data from event repositories and format it for display by the common event console. Each configured connector retrieves event data from a specific event repository.

The status and operational capability of an event connector can be affected by various conditions such as network outages or problems on the event repository server.

You can determine connector status in two ways:

Procedure

- To see an indication of the overall status of all connectors, check the common event connector status icon on the toolbar. The icon indicates the highest-severity current status for any configured connector. The possible status values are as follows, from highest severity to lowest:

Icon	Status	Description
🔵	Unknown	The console is unable to determine the status of any of the event connectors.
🔵	Configuration error	At least one connector is not currently functioning because of an error in the connector configuration.
🔵	Not operational	At least one connector is not currently functioning because the server is not reachable.
🔵	Unstable	At least one connector has experienced errors while retrieving information from the event repository, but it continues to function.

Icon	Status	Description
	Initializing	At least one connector is currently starting up.
	Operational	All connectors are functioning correctly.
	No connectors defined	No common event connectors have been configured.

- To view the status of all configured common event connectors, click the **Common event connector status** icon on the toolbar. The Common Event Connector Status window opens, showing detailed connection status information for each configured connector.

Adding a Tivoli Enterprise Console event viewer

If your environment includes the Tivoli Enterprise Console, you can add it to any workspace to integrate alarms and events from network management operating systems and correlate them with other events.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want the Tivoli Enterprise Console event viewer. The view will show events from the managed systems assigned to this branch of the Navigator. For example, if you add it to a workspace for the Linux Systems Navigator item, all the Linux-based managed systems on this branch of the Navigator will be included in the event reporting.
2. If you want the view to occupy a new space, click **Split Vertically** or **Split Horizontally** in one of the views.
3. Click **Tivoli Enterprise Console**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
4. Click inside the view where you want the Tivoli Enterprise Console event viewer. After you click, the old view is replaced by the Tivoli Enterprise Console event viewer and a login or configuration window might display.
5. If the Tivoli Enterprise Console Login window opens, complete the fields to specify the Tivoli Enterprise Console managed host and your user ID.

Results

After you click **OK**, a dynamic or event group query is issued and events are retrieved from the event server as they occur. The managed systems that are reporting events to this view are those assigned to the current Navigator item.

What to do next

Additional windows for viewing event details and running IBM Tivoli Monitoring tasks can be launched off this view.

When you return to this view during future work sessions, you can log on automatically. If you are logged onto the portal server from your browser, be sure to exit and then restart the browser before attempting to log in again. Otherwise,

any Tivoli Enterprise Console event viewers you add to workspaces will not receive events.

Tivoli Enterprise Console event viewer and console toolbars

The Tivoli Enterprise Console event viewer has tools for working with events and filtering them. It also has a pop-up menu for actions to take on the selected row. When a dialog box is displayed you can click **Help** to access the Tivoli Enterprise Console help system.

Preferences

> This tool opens the Preferences window. The changes you make are for your user ID and affect this and all new Tivoli Enterprise Console event viewers. General preferences set the refresh rate for displaying new events, the age of closed events to display, and the number of events to display. The Working Queue/All Events preferences determine the column sort order in the Working Queue and All Events queue. Severity Colors preferences determine whether or not hover and column highlighting are activated, and what colors to use for event severity.

New Task with Selected Event

> This tool opens a window for configuring automated tasks, which run when a particular event is received by the event console. For example, you can configure an automated task to send an e-mail message to an administrator when an event matching the criteria you defined ahead of time is received by the event console.

> When an automated task completes, it issues information about its success or failure and any results created by the task. An icon is displayed in the Task Completion column of the Event Viewer when an automated task completes. You can double-click this icon to view the task results. This icon also is displayed for rule-based tasks that have completed.

> Automated task results are written to a text file, which is configured in the Console Properties editor for each console.

Reload Events

> This tool is used to clear the view and reload the events. A query is sent to the event database and the results are displayed in the view.

Open Automated Task

> Use this tool to open the Automated Task window for viewing, editing, creating, and deleting automated tasks.

Hide Severity

> These tools show the number of events for each severity and enable you to filter out Unknown, Harmless, Warning, Minor, Critical, or Fatal events or any combination thereof. The example here shows that *Harmless* events have been excluded and the number received is 0. These are toggle tools: click again to resume display of that event severity.

Hide Status

> These tools filter events by their status: Open, Response, Acknowledged, and Closed. These are toggle tools: click again to resume display of events of that status.

Hide Operators' Events

> These tools hide the events for your user ID or those of all other operators signed on to the event server. These are toggle tools: click again to resume display of the events for your user ID or other operators.

Task Completion

An automated task is configured to run when certain types of events have been received by the event console. When an automated task completes, it issues information about its success or failure, and any results created by the task. The task-completion icon is displayed in the first column of the Event Viewer after an automated task or rule-based task has completed. Double-click this icon to view the task results.

Event view rows

Each event received by the Tivoli Enterprise Console adds a row to the view showing summary information. When you select an event (or multiple events with Ctrl+click or Shift+click), you can right-click to get a pop-up menu of options. The buttons are also enabled.

Buttons

Select an event row to enable the buttons to **Acknowledge** the event, **Close** the event, open a tabbed window with event **Details**, or display **Information** about the event. The **Refresh** button shows the event server icon. The button is disabled when no events have arrived since the last refresh of the view; and enabled when new events have arrived.

What to do if the view is disabled

IBM Tivoli Enterprise Console relies on the Tivoli Management Framework for common services, including communications. If you are running Tivoli Enterprise Portal in browser mode and log out of the portal server, either intentionally or unintentionally, contact between IBM Tivoli Enterprise Console and the Tivoli Management Framework is lost.

No new events will be added to the view even if you log out and log in again. You must recycle (close the application, then start it again) the browser before logging in again to the portal server before the IBM Tivoli Enterprise Console event viewer can resume event reporting.

Note: When a situation has been stopped that has an open event that was forwarded to the IBM Tivoli Enterprise Console, the effect is that all open IBM Tivoli Enterprise Console events whose cms_hostname and situation_name (as specified in the event synchronization rules) is the same as that of the stopped situation will be closed. This is done to maintain event status synchronization between the monitoring server and the event server.

Creating a browser view

The browser view invokes the browser that comes with Tivoli Enterprise Portal so that you can open web pages and web applications in the workspace. When saved with the workspace, the browser view opens with the workspace at the URL used when it was saved.

Before you begin

This feature requires your user ID to have Workspace Author Mode permission.

About this task

The Tivoli Enterprise Portal integrated browser uses little memory and provides basic browser functionality. It does not have many of the extended features that come with some external browsers. To add a browser view to the current workspace:

Procedure

1. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
2. Click ▣ **Browser**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
3. Click inside the view where you want the browser. After you click, the old view is replaced by the browser.
4. In the **Location** text box, type the URL of the web page to open.

Browser view toolbar

The browser view contains a toolbar and a text box for entering specific URLs, which you can use to navigate through pages.

⬅ **Back**	Returns to the previous page
➡ **Forward**	Moves forward to the next page in your browsing sequence.
⬤ **Stop**	Stops loading the page.
⟳ **Refresh**	Refreshes the page. If the web page has been recently updated, this tool refreshes the view with the latest version.
⌂ **Home**	Goes to the start page for this workspace. Initially, this is set to the Help overview topic. If you change the URL, then save the workspace, the new URL becomes the start page for this browser view.
🖨 **Print**	Prints the current web page.
🔍 **Find**	Brings up the Find window for searching the web page for the text specified.

URL to files on the portal server

Tivoli Enterprise Portal provides a relative URL reference for HTML files you would like to share. You can place HTML files on the Tivoli Enterprise Portal Server computer in the *install_dir*\cnb (/cw on Linux) folder so that they can be referenced in the browser view and saved with the workspace. Then any user who opens that workspace will see the HTML page in the browser view.

This feature is useful in environments where browser mode users pass through a firewall to log on to Tivoli Enterprise Portal.

Copy the files you want to reference to the *install_dir***cnb** (/**cw** on Linux) folder or subfolder on the computer where the Tivoli Enterprise Portal Server and browser client framework are installed.

In the browser view, enter the file name in the **Location** address field and any path information if it is in a subfolder. For example, the entry would be

information.htm if the file is in the **cnb** or **cw** folder, or my_files/information.htm if the files are in the **my_files** subfolder under **cnb** or **cw**.

After you press Enter to open the HTML page, the address is prefixed with http://*localhost*:15200/ (where *localhost* is the fully qualified name of the computer where the portal server is installed) and the page is opened.

Link symbols in the URL address

You can include symbols in the location URLs that can be substituted with link symbol values stored in the workspace's property context pool.

The symbol for substitution must be of the following format:

$symbol_name$

Before the browser view sends the URL request, it replaces all the symbol names with the associated values from the context pool.

Link to external browser

You can launch an external browser from within a hyperlink in the browser view. This is the mechanism used to view attachments to events.

If you want a link to launch an external browser, use the following attributes in the hyperlink:

TEP_TARGET=external
> If this attribute is found on the <a> link tag, the URL is launched in default external browser.

TEP_LOCATION=relative
> If this attribute is found on the <a> link tag, the "href=" value is treated as a relative URL and the path is appended to the web server's based resources.

Adding a notepad view

The notepad view gives you a text editor for writing about the workspace.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new pane, click ⊞ Split Vertically or ⊟ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
3. Click 📄 **Notepad**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
4. Click inside the view where you want the notepad. After you click, the old view is replaced by the notepad and a cursor is displayed.

What to do next

Type any text in the view. The standard text editing and cursor movement keys on your keyboard can be used with the notepad view.

Backspace
> Delete the character to the left of the cursor.

Delete Delete the character to the right of the cursor.

Arrows
> Move the cursor left or right one space, or up or down one line.

Home Move the cursor to the beginning of the line.

End Move the cursor to the end of the line.

When you are finished, save (or File > Save As) the workspace if you want to keep the notepad view and the text you entered.

Creating a take action view

Tivoli Enterprise Portal enables you to interact directly with your applications and operating system through its take action feature. You can use the take action view to enter a system command or to stop or start a process at the system.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
3. Click ⚡ **Take Action**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
4. Click inside the view where you want the take action view. After you click, the old view is replaced by the take action view.

What to do next

If you want to keep the take action view in this workspace for future work sessions, either click 🖫 **Save** to update the workspace properties with the new view, or click **File** > **Save Workspace As** to save this as a new workspace and leave the original unchanged.

Terminal view

Use the terminal view to connect to a TN3270, TN5250, or VT100 host system with TCP/IP. For a 3270 or 5250 terminal view, you also have scripting capability with record, playback, and authoring of entire scripts.

By associating a terminal view with a connection script and a query that returns appropriate values, you can configure a view that opens to a specific panel of a

3270 application. This feature, called *dynamic terminal integration support*, is useful for creating contextual workspace links for investigating issues.

Your user ID must have Workspace Author Mode permission to use this function.

Adding a terminal view to the workspace

You can add a terminal view to your workspace to connect to another system. This process is also known as Host-On-Demand. For example, you can connect to a TN3270, TN5250, or VT100 host system with TCP/IP, and then save the connection settings with the workspace.

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new pane, click ⊞ Split Vertically or ⊟ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
3. Click ▣ **Terminal**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
4. Click inside the view where you want the terminal emulator. When you click inside a view, the Terminal Emulator Configuration window opens.
5. Specify the connection information for the view, or specify that you want to retrieve the connection information from a query.
   - If you want to specify the connection information in the configuration, select **Use the values from this dialog**, and fill out the appropriate fields:

     Host Type the host name or IP address of the system to which you want to connect.

     Port Type the IP port number to which you want to connect. The default of 23 is valid for most installations.

     Terminal type
     Select a terminal type from the ▾ Terminal Type list or accept the default.
   - If you want to retrieve connection information from a query, select **Fetch values from a query**. You will then need to specify the appropriate query in the view properties.

     Note: To support obtaining connection information from a query, the monitoring agent for the target application must provide a suitable query to use. For more information, refer to the agent documentation.
6. Click **OK** to accept the configuration and add the view.
7. If you specified **Fetch values from a query**, open the view properties and specify the query that will provide the connection information. To determine which query to use, refer to the agent documentation.
8. If you want to keep the terminal view in this workspace for future work sessions, do one of the following:
   - Click 💾 **Save** to update the workspace properties with the new view.
   - Click **Save Workspace As** in the **File** menu to save this as a new workspace and leave the original workspace as it was.

Results

This configuration information will be used each time the workspace is opened and will establish a connection to the specified host. By default the terminal connection remains active until you disconnect manually or until you end this Tivoli Enterprise Portal work session. This means you can open other workspaces and return to this one without breaking the connection.

What to do next

If you want to use dynamic terminal integration support to link to a specific application panel, you must also configure a connection script. Without a connection script, users will need to manually navigate to the appropriate application panel.

What the terminal view shows

The terminal view displays a toolbar, main terminal window, keyboard mapping, and status bar.

About this task

The terminal view has four parts:
* A toolbar to control the display and operation of the session:

  KeyPad toggles the display of the key mappings at the bottom of the view.

  Connection connects to a session or disconnects from an active session. Disconnecting from a session while a script is active will terminate the script; disconnecting while the script is recording will stop the recording.

  Manage Scripts opens a window so you can create, edit, delete, start, and stop scripts for 3270 or 5250 sessions. Scripts are not available   for Telnet sessions.

  Record New Script prompts for a new script name and starts recording. The text box shows `Recording:` *name*, where *name* is the title of the script currently being recorded. Scripts are not available for Telnet sessions.

  Stop Script stops whichever is in use: script recording or script playback.
* A main terminal window that represents the emulator session with the host. You manipulate this window in the same manner as a normal terminal.
* A keyboard mapping that emulates the aid keys of the real terminal. You can click directly on these keys to send the equivalent aid key.
* A status bar that displays information about the current status of the session:

  Host name displays the name of the host to which you are connected.

  Connection status indicates whether the current session is connected or disconnected.

  Keyboard status indicates whether the keyboard is locked or unlocked.

  Overwrite/Insert indicates whether the characters will overwrite or insert when typed into an unprotected field.

  Field Status indicates whether a field on the terminal screen is protected or unprotected.

  Cursor Location displays the screen row, column of the cursor.

Terminal view properties

The terminal view has Configuration, Style, and Query properties that can be edited to suit your environment.

Click ✎ **Edit properties** in the terminal view title bar or right-click the view and click ▦ **Properties** to see these tabs:

Configuration

The ▤ **Configuration** tab has sub-tabs for customizing the terminal view and its display: **Connection** and **Scripts** and **Zoom**. Whenever you open the workspace, the script you specify will run automatically and the host you specify will be connected automatically.

⚲ Connection

Connection

> This is the same Terminal Emulator Configuration window you get when adding a terminal view to the workspace. The values are those you supplied when you added the terminal view. You can change them here:
>
> ⦿ **Fetch values from a query** indicates that you want to retrieve connection information from a query. You must specify the appropriate query on the Query tab.
>
> ○ **Use the values from this dialog** indicates that you want to use the connection information specified in these fields: **Host** is the host name or IP address of the system to which you want to connect; and **Port** is the IP port number used by the host software. The default of 23 is valid for most installations.
>
> **Terminal Type** is the name and dimensions of the terminal to emulate. Accept the default or select one from the ▾ list.
>
> ☑ **Connect when entering workspace** will connect to the host when you open the workspace. If you disable it, you must click ⚲ Connection to start the session.

Disconnect options

> ⦿ **Do not disconnect when leaving workspace** keeps the session open as long as you are logged on to Tivoli Enterprise Portal, so you can return to this workspace and the session where you left it last.
>
> ○ **Disconnect after this many minutes:** ▭ will disconnect the works session the number of minutes entered here after you close the workspace.

🗋 Scripts

This is the same as the Manage Scripts window, which opens when you click ▤ **Manage Scripts** in the terminal view.

Style

The 🇦 **Style** tab lets you specify the font and font size used for the session and several terminal display options.

Text style

> **Font name** for the session display is Monospaced by default or you can select another from the list.

Size is 12 by default or you can select another point size from the list.

Terminal display options

☐ **Show keypad** selects ▦ **KeyPad** automatically for the view whenever you open the workspace, so the function key mappings will display at the bottom of the view.

☐ **Show block cursor** displays the cursor as a block instead of as the underscore character.

☐ **Show ruler** draws a line along the vertical axis of the cursor and along the horizontal axis so you can easily locate the cursor on the display.

Query

Use the Query tab to specify the query to use for the Terminal view. The values returned by the query can be used in two ways:

- If you specified **Fetch values from a query** on the Connection tab, the query provides the host name or IP address, port, and terminal type to use for the connection.
- The query provides application-specific values that can be retrieved using the getProperty function in a connection script. These property values provide contextual data that can be used (along with any properties specified by the link) to open the Terminal view to a specific panel.

If you are using a query to provide connection information, the query you select must return a single row of data that includes certain fields. These are the minimum required fields for Terminal view query:

Attribute	Description
Origin_Node	Origin node
HOSTIPADDR	Host name or numeric IP address to connect to
TN3270PORT	TCP/IP port to connect to
LUGROUP	LU group
EMU_TRMTY	One of the following terminal types: • IBM VT100 (24x80) • IBM 3270 (24x80) • IBM 3270 (32x80) • IBM 3270 (43x80) • IBM 3270 (27x132) • IBM 5250 (24x80) • IBM 5250 (27x132) This field is optional; the default terminal type is 3270 (24x80).

In addition to the minimum required fields, the query can also return additional fields to supply application-specific data. The connection script can access these fields using the getProperty function.

Running a script

After you have scripts saved for a terminal view, and the terminal view saved with a workspace, you can run the scripts at any time. You must be connected to a 3270 or 5250 session before you can run a script.

About this task

To run a script:

Procedure

1. While in the terminal view, click 🐾 **Connection** and fill in the Terminal Emulator Configuration window. You can specify the connection information in the window, or select **Fetch values from a query** to use the connection information returned by the query specify in the view properties. Click **OK** to accept the configuration and establish the connection.
2. Click 📋 **Manage Scripts** to open a list of available scripts.
3. Select a script from the list or type a name in the **Script Names** field.
4. Click **Run** to start executing the selected script.
5. If you want to stop the script before its completion, click ⬤ **Stop**.

Recording a script

The recorder allows you to capture and play back host 3270 or 5250 sessions. As you interact with the terminal session, the session is recorded as a set of script commands that can be saved and played back at a later time. This allows for the automation of navigation to specific screens.

About this task

To record a script:

Procedure

1. While in the terminal view, click 👥 **Record New Script**.
2. Type a new *name* in the **Script Name** field and click **OK**. From this point on, all terminal activity will be captured in the script.
3. Perform all the actions to be recorded in the script.
4. When you are finished, click ⬤ **Stop**. The script is now available in the workspace for editing or execution through 📋 **Manage Scripts**.
5. If you want the script to start automatically upon connection, click 📋 **Manage Scripts** and select the script from the Startup Script ▾ list.
6. Click 💾 **Save** to permanently save the script. Session scripts are saved with your user ID, so no other users will be able to see and run the same scripts unless you have Workspace administration mode turned on when you create them. The saved script is available from any workspace to which you add a terminal view.

Managing scripts

Use the Manage Scripts function to edit existing scripts and to specify a script to run whenever you start a 3270 or 5250 terminal session.

About this task

To manage scripts:

Procedure

1. While in the terminal view, click 📄 **Manage Scripts** to see the list of available scripts. There are two types of scripts:

   - **Global scripts** are stored as part of the Tivoli Enterprise Portal configuration and are available in any workspace. In the **Manage Scripts** window, global scripts appear in normal text.

   - **Local scripts** are stored as part of the workspace configuration and are associated with a particular agent. In the **Manage Scripts** window, local scripts are indented and appear in italic text.

2. Select a script from the list, or type a name in the **Script Names** field, then click one of the following:

   - **Run** to start executing the selected script. Note that you must be connected to a session before you can run a script.

   - **Stop** to discontinue running the script. This is useful if you want to stop a script before it has completed or if an unrecoverable error has occurred.

   - **Record** to start recording a new script. The Record New Script window appears, prompting you for details about the new script.

     Note: Select the **Local script** check box to specify that the new script should be saved as local rather than global. Saving scripts as local is recommended because local scripts can be exported and restored as part of the workspace configuration. (For more information, refer to the *IBM Tivoli Monitoring Administrator's Guide*.)

   - **Stop Recording** to stop the recorder.

   - **Open** to edit the selected script. The Script editor opens with basic editing functions such as Copy and Paste. It has menu items to create a **New** script, **Save** the script, **Save As** another script, and **Exit** the Script editor.

   - **Delete** to remove the selected script from the list of available scripts.

     Note: Some scripts are read-only and cannot be deleted or modified. If you edit a read-only script, you can save the modified version as a new script by clicking the **Save As** menu choice.

3. If you want to designate a script to start automatically upon connection, select if from the Startup Script ▾ list. Connection to a host system is established whenever you open a workspace that contains a terminal view, and when you reconnect to a session with 🖇 **Connection** or by running a script. Instead of using hard-coded values, the connection script can use the getPropery function to retrieve values specified by a workspace link or by the query specified in the view properties. These values can include connection information as well as application-specific data used to navigate to a specific application panel for investigating issues.

4. Click 💾 **Save** if you designated a startup script.

Terminal Emulator Scripting Language

Automate terminal view session navigation with scripts written with our proprietary Terminal Emulator Scripting Language and saved for use in your monitored environment.

The scripting engine is an interpreter/execution engine based on a C-like language. It has a defined format and syntax, and a set of built-in functions. It also allows you to create your own functions. Because it is intended to address a specific

environment—that of manipulating terminal session navigation—it does not provide a full range of programming functionality.

Script syntax

Create new 3270 or 5250 session scripts using the Manage Scripts feature that is available from the terminal view.

A script is a program made up of function definitions, variable definitions, and executable program statements. The terminal emulator scripting language is similar to Java or C. The scripting language has these general conventions:

- Keywords, variables and function names are case-insensitive
- All statements are terminated with a semi-colon
- Multiple statements can appear on a single line
- Function definitions must appear ahead of the data declarations and program statements
- Declarations can appear anywhere before they are referenced

The general layout is:

Return-type FUNCTION function-name (optional parameters) {optional local variables; statements...}

Reserved words

Reserved words are character strings that have special meanings. They cannot be used except for their intended meaning. The reserved words are of three types:

Language keywords

Language keywords are used as commands within a script. They are interpreted to provide some action or considered as part of a statement. You can use these keywords only for their predefined purpose:

AND	; (semicolon)	If
OR	Break	Int
EQ	Continue	Real
NE	Else	Return
GE	Exit	String
GT	For	While
LE	Function	LT

Declarations and function names

Each script program begins with global data declarations, stating the global variables and user-defined functions you will be using. Declarations build an association between a function and its attributes or values. You do not need to declare any of the built-in functions, because the interpreter already recognizes these function names.

Identifiers

Identifiers are the names that you create to denote constants, variables, and functions. Identifiers have the following characteristics:

- an identifier can be any length
- the first character must be alphabetic (a-z, A-Z), numeric (0-9), or an underscore (_)
- the remaining characters can be alphabetic (a-z, A-Z), numeric (0-9), or underscores (_)

- each identifier must be unique. Identifiers cannot be one of the reserved words.

Punctuation

The following punctuation rules apply to a section

- Statements end with a semi-colon – ;
- Parameter lists are enclosed in parenthesis – ()
- Parameters are separated by commas – ,
- Statement blocks are enclosed in braces – {}

Embedded quoted strings

String literals are enclosed in double quotation marks or single quotes. You can insert double quotation marks as part of a string literal by preceding the quote with the escape character /". Alternatively, you can begin and end the string literal with the single quote character '. To imbed a single quote use the /' escape sequence or enclose the string literal within double quotation marks. Example:

```
// Both statements yield embedded quotes quotedStr1 =
"The job /"TSOA/" is running"; quotedStr2 = 'The job "TSOA" is running';
```

White space

The script language ignores white space (spaces, tabs, new-lines) except within a quoted string.

Comments

Comments are supported as a way to add explanatory text to your script program or to exclude certain parts of the code.

- comments can be single line or multi-line comment blocks
- the script interpreter ignores comments
- comment blocks are enclosed within the /* */ pairs. Comment blocks can span multiple lines.
- comments can also start with the // character string. The interpreter ignores everything to the right of the double slash up to the end of the line (new-line character).
- comments can start anywhere in the script

Example:

```
// This is a single line comment a = 3;
// this is a comment /* This section defines multiple line
comment block */ a = 3; /* imbedded comment block */ b=4;
```

Constants and variables

A constant is a named data item with a predefined value, while a variable is a named data item whose value can change during the course of program execution. Variables can be either global or local.

All data is represented as one of three basic data types:

Int   Integer data type (+ - followed by 1 or more digits 0–9).

Real   Floating point (+- followed by zero or more digits 0–9). Decimal point followed by 1 or more digits 0–9. Note that all numbers are stored internally as Float.

String Any valid Unicode string enclosed within quotes.

Constants

A constant is a named data item with a predefined value. You cannot change the value assigned to a predefined constant. The predefined constants are:

NULL An empty reference. Similar to an empty pointer. Note that this is not the same as a null string "".

TRUE

Equivalent to the number 1.

FALSE

Equivalent to the number 0.

Variable definitions

A variable is a named data item whose value can change during the course of program execution. A variable name must follow the naming convention of an Identifier (alphabetic character, numeric character, and the underscore.) When more than one variable is defined on a single declaration, the name must be separated by commas. Each variable declaration must be terminated with a semi-colon. Variable names cannot be the same as a reserved word.

Note: Identifiers and Names in the script language are case-insensitive. Thus, abc is the same as ABC.

Global and local variables

Variables can be either global or local. A variable is global unless it is declared within a function definition. Global variables are visible and available to all statements in a script. Local variables are only visible and available within the function where they were defined. Although variable names and identifiers must be unique, it is valid to declare a local variable that has the same name as global, or the same name as a local variable defined in another function. Parameter names in a function definition are considered to be local variables. Data passed to a function is by value. A global variable passed to a function will not be changed by the function.

Operators and expressions

Expressions perform specific actions, based on an operator, with one or two operands. An operand can be a constant, a variable or a function result. Operators are arithmetic, logical, and relational. As with C, some operators vary in functionality according to the data type of the operands specified in the expression.

Arithmetic operators (+, -, *, /, **, %)

Arithmetic operators perform mathematical operations such as addition and subtraction with operands. There are two types of mathematical operators: unary and binary. Unary operators perform an action with a single operand. Binary operators perform actions with two operands. In a complex expression, (two or more operands) the order of evaluation depends on precedence rules.

Unary arithmetic operators

Unary operators are arithmetic operators that perform an action on a single operand. The script language recognizes the unary operator negative (-).

The negative unary operator reverses the sign of an expression from positive to negative or vice-versa. The net effect is that of multiplying the number by -1. Example:

```
a = -10;
```

The Prefix operators increment or decrement the variable prior to dereferencing the object, while the Postfix operators increment or decrement the variable after referencing it. Example:

```
A=1;
B = a++; // b will equal 1, a will equal 2;
A = 1;
B = ++a; // b will equal 2, a will equal 2;
A= 1;
B = a--; // b will equal 1, a will equal 0;
```

Binary arithmetic operators

Insert a space before and after an arithmetic operator. The binary arithmetic operators that are supported are listed below.

Table 1. Binary Arithmetic Operators

Symbol	Operation	Example	Description
+	Addition	a + b	Add the two operands
-	Subtraction	a - b	Subtract the second operand from the first operand
*	Multiplication	a * b	Multiply the two operands
/	Division	a / b	Divide the first operand by the second operand
**	Power	a ** b	Raise the first operand by the power of the second operand
%	Modulo	a % b	Divide the first operand by the second operand and yield the remainder portion

Operator precedence

Expressions are normally evaluated left to right. Complex expressions are evaluated one at a time. The order in which the expressions are evaluated is determined by the precedence of the operators used. The standard C ordering is followed.

1. negation (-) unary
2. power
3. multiplication, division and modulo
4. addition and subtraction

If an expression contains two or more operators with the same precedence, the operator to the left is evaluated first. For example, 10 / 2 * 5 will be evaluated as (10 / 2) and the result multiplied by 5.

When a lower precedence operation should be processed first, it should be enclosed within parentheses. For example, 30 / 2 + 8. This is normally evaluated as 30 divided by 2 then 8 added to the result. If you want to divide by 2 + 8, it should be written as 30 / (2 + 8).

Parentheses can be nested within expressions. Innermost parenthetical expressions are evaluated first.

Assignment Operator (=)

Use the assignment operator (=) to copy a constant, literal, variable expression result, or function result to a variable. The script language does not support multiple assignments in a single statement (such as a=b=c=0). String lengths are defined based on the size of the string assigned to the variable and can change dynamically at runtime.

Logical Operators (AND, OR)

Logical operators allow the combining of more than one relational test in one comparison. Logical operators return a TRUE (1) or FALSE (0) value. Logical operators have a lower precedence than arithmetic operators.

Table 2. Logical Operators

Symbol	Operation	Example	Description
AND &&	AND	Expr1 $$ expr2	True if both expr1 and expr2 are true.
OR ||	OR	Expr1 OR expr2	True if either expr1 or expr2 are true.

Relational Operators

Relational operators are as follows:

Table 3. Relational Operators

Symbol	Operation	Example	Description
<	Less than	a < b	True if a is less than b.
>	Greater than	a GT b	True if a is greater than b.
==	Equal	a == b	True if a is equal to b.
!=	Not equal	a NE b	True if a is not equal to b.
<=	Less than or equal	a <= b	True if a is less than or equal to b.
>=	Greater than or equal	a GE b	True if a is greater than or equal to b.

Comparisons must be made on like data types—string to string, numbers to numbers. If the data types are different, a run time error will be raised.

String comparisons

Two strings are considered equal if they match character for character and are of the same length. Strings are compared character by character, left to right until all characters have been matched or one of the strings has been exhausted. If the end of one string is encountered before the end of the corresponding string, the strings are considered to be of unequal length and result in an unequal condition.

If any character in a string does not match the corresponding character in the other string, the comparison stops and the strings are considered not equal. Strings are case-sensitive. Examples:

```
Str1 = "abcdefg";
Str2 = "abcdefg";
Str3 = "abcdef";
Str4 = "ABCDEFG";
If (str1 == str2)... results in an equal/true conditions
If (str1 == str3).... Results in a not equal condition
because str3 is shorter than str1
If (str1 == str4) ... Results in a not equal condition
because the strings are of different case;
```

String Concatenation Operator (+)

The Plus(+) operator, when applied to strings, results in the concatenation of the two strings. Example:

```
Str1 = "abc";
Str2 = "def";
Str3 = str1 + str2; results in "abcdef"
```

Flow control statements

The script interpreter processes statements sequentially starting with the first statement of the program block. This sequential flow can be controlled with conditional statements that perform branching and iteration. Branching is controlled with an IF statement that directs execution along a given path. Iteration is controlled with loop statements that execute one or more statements repeatedly. Recursion is the process of a function calling itself either directly or indirectly through a second function.

Control statements are:

IF/THEN/ELSE, WHILE/BREAK/CONTINUE, RETURN

IF statement

The **IF** statement allows conditional execution based on the evaluation of an expression. **IF** statements can be nested.

Syntax:

IF (condition) true-statements [ELSE false-statements] ;

Table 4. If Statement

Part	Description
IF	The keyword to begin the control structure.
Condition	A numeric or string expression that evaluates true (one) or false (zero).
True-statements	Statement(s) performed when the condition is true.
ELSE	(Optional) Keyword to begin the False control flow.
False-statements	Statement(s) performed if the condition is false.

Multiple statements are permitted as part of the true-statements or false-statements. If multiple statements are to be executed, they must appear within a statement block embedded within braces { }. Example:

```
If (userid != "") setField(userid,1,5);
If (userid != "") setField (userid,1,5)
else
errMsg("Userid must be specified");
If ( op == 5 )
{ setField(userid,1,5);
sendKey("Enter");
}
else
if (op == 3)
setField(pswd,2,5);
else
msgBox("unknown operation:);
if (prompt("enter Name") == "") errMsg("Name must be specified");
```

For statement

The **For** statement allows iterative execution based on the evaluation of an expression. An initializer statement can be executed to establish initial conditions. A conditional expression is used to control the number of times the loop is executed. An update statement is issued each iteration through the loop.

Syntax:

FOR (initial-statement ; condition ; update-statement) loop-statements ;

Table 5. For Statement

Part	Description
For	The keyword to begin the loop control structure.
Initial-statement	A statement to be executed prior to the loop.
Condition	A numeric or string expression that evaluates true (one) or false (zero). A function that returns an integer result. The loop continues to execute while this condition is true.
Update-statement	A statement that will be executed at the end of each iteration.
Loop-statements	Statement(s) performed when the condition is true.

Multiple statements are permitted as part of the loop-statements. If multiple statements are to be executed, they must appear within a statement block embedded in braces { }. Example:

```
For ( i=0; i<3; i=I+1) waitforscreen(1000);
For ( i=0; i<3; i=i+1)
{
sendString(userid);
if (waitforstring ("Logon complete") > 0)
{
success = true;
break;;
}
}
```

Return statement

Use the **Return** statement to return from a function or to end the main program block. Program execution resumes at the statement following the point at which a function was called, or exits the script if it is in the main program block. An optional value can be associated with the return statement.

A RETURN statement that is not within a function definition terminates the script.

Syntax:

RETURN [return-value];

Table 6. Return Statement

Part	Description
RETURN	The keyword to pass control to the calling statements.
return-value	An optional value to be returned to the caller. This can be a constant value, literal, or expression.

Example:

```
.....
return 0;
```

While statement

The **While** statement executes a series of statements in a loop for as long as a condition is true.

Syntax:

WHILE (expression) statements;

Table 7. While Statement

Part	Description
While	The keyword to begin the control structure.
Condition	A numeric or string expression that evaluates true (one) or false (zero).
statements	Statement(s) performed while the condition true.

Example:

```
cnt = 1;
WHILE (cnt < 3)
{
a = a+1;
cnt = cnt + 1;
}
// this will execute the loop 3 times
```

Break statement

The **Break** statement causes control to exit a While loop regardless of the condition controlling the loop.

Syntax:

BREAK;

Example:

```
cnt= 1;
WHILE (cnt < 3)
{
userid = prompt("Enter Userid");
if (logon(userid) == 1) break;
cnt = cnt + 1;
}
```

Continue statement

The **Continue** statement causes the execution of the While loop to resume at the top, without executing the rest of the statements in the block.

Syntax:

Continue;

Example:

```
cnt= 1;
WHILE (cnt < 3)
{
userid = prompt("Enter Userid");
if (logon(prompt("Enter Userid")) != 1)
{
msgBox("Invalid Logon id");
continue;
};
// do other logic....
}
```

User-defined functions

User-defined functions make it possible to create structured scripts. User- defined functions must be declared at the beginning of a script. The FUNCTION keyword identifies the block of statements as a user-defined function.

The function declarations must follow a specific format:

return-type FUNCTION function-name (optional parameters) { optional local variables; statements...}

Part	Description
Return-type	The data-type of the return value of the function.
FUNCTION	The keyword that identifies this as a user-defined function.
Function-name	A unique name that identifies this function. Must follow the naming convention of Identifiers.
Optional parameters	The named parameters (if any) that will be passed to this function. Enclosed in parentheses.
Optional local variables	The declaration of any local variables that this function will use.
statements	The Statement(s) that will be performed when the function is called.

The rules described in the "Constants and variables" topic apply to the parameters and local variables used in a function. Any optional parameter names are assumed to be local to a function. Other local variables that are unique to a function can be declared. The control flow of the statements in a function follow the same rules that apply to normal program statement blocks. A RETURN statement returns control to the point in the program where a function was called. A return value might be passed back as part of the return statement. Example:

```
int FUNCTION signonPanel(string arg1)
{
if ((findString(arg1) > 0) AND
(findString("Sign On Panel") > 0) AND
(findString("Userid") > 0) AND
(findString("Password") > 0) )
return 0;
else
return 1;
}
```

Function calls

A function is a uniquely named set of statements that perform a certain task. A function can take a set of arguments to operate on and return a value as a result of the task it performed. Use a function's value in assignment statements, condition statements or as a program statement.

All functions, whether user-defined or built-in, are called in the same manner: enter the function name followed by the parameters to be passed to the function. Example:

```
signonPanel("OMEGAVIEW");
```

A function's value can be used in assignment statements, condition statements or as a program statement. Examples:

```
Result = signonPanel // an assignment
("OMEGAVIEW");
If (signonPanel("OMEGAVIEW") == // in a conditional statement;
1)
SignonPanel("OMEGAVIEW"); // as a program statement
```

Built-in functions

The terminal script language has a set of built-in functions that can be used to manipulate data as well as to control the navigation through a terminal session. Because these functions are already defined internally, they do not have to be declared before using them. Simply type the name of the function.

The following sections describe the built-in functions and their parameters and return values.

Upper Upper converts a string to upper case.

> **Parms** String to be converted
>
> **Return**
>> Returns a new string, converted to uppercase. Follows the rules of the language in use.
>
> Example:
> ```
> Str = "abc";
> Str2 = upper(str); // str2 will contain "ABC"
> ```

Lower Lower converts a string to lower case.

> **Parms** String to be converted
>
> **Return**
>> Returns a new string, converted to lowercase. Follows the rules of the language in use.
>
> Example:
> ```
> Str = "ABC";
> Str2 = lower(str); // str2 will contain "abc"
> ```

Index Index returns the position (index relative to one) of a substring within a String.

> **Parms** Target string - string to be searched
>
> **Return**
>> Returns the position of the search string or zero if not found
>
> Example:
> ```
> Str = "abcdef";
> Val = index(str,"def"); // val will contain 4
> ```

Length

> Length returns the length of the string.
>
> **Parms** String to be inspected
>
> **Return**
>> Returns a numeric value of the length of the string (in characters), or 0 if null string
>
> Example:
> ```
> Str = "abc";
> Val = length(str); // val will be 3
> ```

Substring

> Substring returns the portion of the string as specified by the offset and length. It can be abbreviated as Substr.
>
> **Parms** Target string - string to be searched
>
>> Starting offset - offset within the string to start the selection Length - length of the data desired
>
> **Return**
>> Returns the string that represents the selection

Note: The length parameter is optional, and if not specified will assume the length up to the end of the target string. If the length exceeds the length of the target string, an error is raised.

Example:

```
Str = "abcdef";
Str2 = substring(str,3,3); // str2 will contain "def"
Str3 = substr(str,1); // str3 will contain "bcdef"
```

Min   Min compares two numbers and returns the value equal to the lesser of the two numbers.

Parms   Number 1 - the first value to compare

Number 2 - the second value to compare

Return

Returns either Number 1 or Number 2 depending upon which is less

Example: Val1 = 10; Val2 = 20; Val3 = min(val1,val2); // val3 will contain 10;

Max   Max compares two numbers and returns the value equal to the greater of the two numbers.

Parms   Number 1 - the first value to compare

Number 2 - the second value to compare

Return

Returns either Number 1 or Number 2 depending upon which is greater

Example:

```
Val1 = 10;
Val2 = 20;
Val3 = max(val1,val2); // val3 will contain 20;
```

Prompt

Prompt allows the script to gather data from a user interactively by asking a user to enter data. This presents a window with message text and an edit field. It prompts you to enter data. You can either enter data and click **OK** or click **Cancel** to ignore the prompt.

Parms   MsgText - text of the message to be displayed in the window

PSWD - optional second parm. The keyword PSWD indicates that this is a password type field. The data entered by you does not display, but is represented by a series of asterisks(*) for each character typed.

Return

Returns the string value of the data entered. If you click **Cancel**, a null value will be returned. This is different from a null string.

Example:

```
Userid = prompt("Enter your userid")
Pswd = prompt("Enter password","PSWD");
```

If you enter a value, it is placed in the variable userid. If you do not enter any data, and click **OK**, a null (zero length) string is returned.

If you click **Cancel**, a null pointer value is returned.

MsgBox

MsgBox displays a text message on the screen. A message severity type can be associated with the message. This causes the system-defined messageType icons to be displayed. You can control the type of buttons to appear with the OptionType parameter.

Parms arg 1 = message string

arg 2 = msg type - a numeric value indicating the type of ICON to display

```
 0 = ERROR_MESSAGE;
 1 = INFORMATION_MESSAGE;
 2 = WARNING_MESSAGE;
 3 = QUESTION_MESSAGE;
-1 = PLAIN_MESSAGE (Default)
```

arg 3 = optionType - a numeric value indicating the types of buttons

```
 0 = YES_NO buttons;
 1 = YES_NO_CANCEL buttons;
 2 = OK_CANCEL buttons;
-1 = DEFAULT ok /cancel button(default);
```

@return = a numeric value representing the button that was pressed to dismiss the window

```
 0 = YES button or OK button;
 1 = NO button;
 2 = CANCEL button;
-1 = CLOSED. The dialog was dismissed by closing
     the dialog window - no button was pressed.
-2 = The message box was not responded to and the
     timeout value was exceeded
```

Example:
```
msgBox("The system is running"); // this will display a plain message,
with OK cancel button
Val = msgBox("Invalid signon. Retry?",0,1); // this will present an
Error type message, with the YES,NO,CANCEL buttons;
```

Note: Because the msgBox is waiting for user input, it could take a long time before the user responds. If it is running unattended, the user will never respond, and the script will stall. For this reason, there is a timeout value associated with this function. If the user does not respond within 10 minutes, the dialog is automatically dismissed and a value of -2 returned.

Sleep Sleep suspends the execution for the time interval specified

Parms Time interval1 - the number of milliseconds to wait

Return

Returns zero if successful; minus one if an error

Example:
```
Sleep(1000); // sleep for 1 second;
Sleep(1000*60); // sleep for 1 minute
```

Note: The maximum amount of time permitted is 15 minutes (1000 * 60 * 15). If it is necessary to wait longer, multiple sleep() functions can be issued.

Note: The sleep is an Active wait. The actual internal value is a 1 second wait. The implementation uses a loop counter to issue the 1 second sleep. This prevents a stall condition and allows the script to be stopped or cancelled.

Emulator functions

The emulator built-in functions perform actions related to the navigation and display of data in the terminal session. The session must be connected to a host in order for these functions to execute. If the session is not connected, the function will return with a default value of zero for numeric return values or a null string for string values. If the function fails, a negative one (-1) will be returned.

Locations on the screen are referenced by Row/Column or by buffer offset (relative 0). Thus, the top left position would be Row 1, Column 1 or Location 0. Functions that take a screen position accept either form. Screen locations (buffer offsets) yield different Row/Column values depending on the screen size. For example, if the terminal has 24 rows and 80 bytes, location 80 will be row 2, col 1. However, if it is a model 5, location 80 will be row 1, col 81.

Disconnect

Disconnects the current session from the host. The script terminates when the session disconnects.

Parms None

Return

Returns 0 if success; -1 if an error

Example:

```
disconnect a ( ); // disconnects the current session
```

FindString

Use FindSring to find the string specified anywhere on the screen.

Parms Arg1 - string to search for

Return

Returns the location of the string, or -1 if the string was not found or the session gets disconnected

Example:

```
user = findString("Userid"); // find the string "Userid" and return its
location if ((findString("OMEGAVIEW") > 0) // is this the correct screen?
```

GetCursor

GetCursor returns the current location of the cursor. This is the buffer offset.

Parms None

Return

Returns the relative location of the cursor, or -1 if an error occurs or the terminal gets disconnected

Example:

```
Pos = getCursor(); // returns the cursor position
```

GetCursorCol

GetCursorCol returns the current column number of the cursor.

Parms None

Return

> Returns the integer value of the cursor column number, or -1 if an error occurs or the terminal is disconnected

Example:

```
Col = getCursorCol(); // get cursor column number
```

GetCursorRow

GetCursorRow returns the current row number of the cursor.

Parms None

Return

> Returns the integer value of the cursor row number, or -1 if an error occurs or the terminal is disconnected

Example:

```
Row = getCursorRow(); // get cursor row number
```

GetDefaultTimeout

GetDefaultTimeout returns the value of the WaitForScreen default Timeout.

Parms None

Return

> Returns the value of the WaitForScreen default Timeout in milliseconds

Example:

```
TimeoutValue = getDefaultTimeout();
// returns the cursor wait for screen timeout value
```

GetField

GetField retrieves the string from the field pointed to by location. The number of characters is determined by the length of the field.

Parms Position - the position of the field on the screen. This can be a ROW,COL or the relative location on the screen.

Return

> Returns the string of characters from the field, or a null string if it is an invalid field position

Example:

```
Applid = getField(5,1); // get the data in the field at Row 5, column 1
Applid = getField(80); // get the data from the field at
location 80 - offset 80 (row 2 col 1 on a model 2)
```

Note: The location can be any valid location on the screen. The field containing that address is referenced. Protected or unprotected data can be retrieved. If the location is invalid, an error is generated and a null string is returned.

GetFontSize

GetFontSize gets the current Emulator screen font size.

Parms None

Return

> Returns an integer value of the Font point size

Example:

```
FontSize = getFontSize();// return the current Font size (in points)
```

GetProperty

GetProperty retrieves the value of a System Property.

Parms PropertyName - the name of the system property to retrieve, enclosed in quotes

Return

Returns the value of the property, or null if not found

Example:

```
Applid = getProperty("APPLID") // get the data from the APPLID variable
```

GetString

GetString retrieves the string from the location on the screen. The number of characters is determined by the length parm.

Parms Position - the position on the screen. This can be a ROW,COL or the relative location on the screen.

Length - number of characters to retrieve

Return

Returns the string of characters from the location

Example:

```
Applid = getString(5,1,16); // get 16 characters from screen
location Row 5, column 1   Applid = getString(80,15);
// get 15 characters from location 80 - offset 80 (row 2 col 1 on a model 2)
```

Note: The location can be any valid location on the screen, protected or unprotected. If the location is invalid, an error is generated and a null string is returned.

Hide Hides the terminal screen. This sets the terminal frame to not visible. It is used when it is not desirable to show the actual screens that are being navigated. When the script ends, or the session is disconnected, the screen will be shown.

Parms None

Return

Returns 0 if successful, or -1 if there is an error or the terminal gets disconnected

Example:

```
Hide(); // set the terminal window to not visible
```

Screencols

Screencols returns the number of columns on the screen. For a 3270 model 2 this would be 80.

Parms None

Return

Returns the number of columns on the screen, or -1 if there is an error or the terminal gets disconnected

Example:

```
Cols = screenCols(); // returns the number of columns
```

Screenrows

Screenrows returns the number of rows on the screen. For a 3270 model 2 this would be 24.

Parms None

Return

> Returns the number of rows on the screen, or -1 if there is an error or the terminal gets disconnected

Example:

```
Rows = ScreenRows(); // returns the number of rows
```

SendKey

> SendKey simulates pressing the corresponding AID key.

> **Parms** String representation of the key to send

> **Return**

>> Returns 0 if successful, or -1 if there is an error or the terminal gets disconnected

> Example:

> ```
> SendKey ("Enter"); // sends the Enter key. Any modified
> data on the screen will be sent as part of the data stream.
> ```

> The following values can be specified:

> **ENTER**

>> Enter key. Sends the modified data, cursor location and the AID key.

> **RESET**

>> Resets and unlocks the keyboard. No data is sent to the host.

> **CLEAR**

>> Sends the CLEAR Aid key. Only the Aid key is sent.

> **SYS_REQ**

>> Sends the System Request Aid.

> **ERASE_EOF**

>> Clears the Unprotected fields from the current cursor location to the end of screen. No data is transmitted to the host.The fields are set to nulls.

> **ERASE_INPUT**

>> Clears the current field to nulls.

> **FIELD_MARK**

>> Transmits the FieldMark aid key.

> **HOME**

>> Puts the cursor at the first unprotected field on the screen.

> **TAB** Places the cursor at the next unprotected field.

> **PA1, PA2, PA3**

> **F1 through F24**

> The following special keys position the cursor:

> **HOME**

>> Puts the cursor at the first unprotected field on the screen.

> **TAB** Places the cursor at the next unprotected field.

> **Note:** The behavior and data that gets transmitted to the host are dependent on the actual key being sent. The Enter and the Function keys

send any modified data fields, as well as the cursor location and the AID key. The other AID keys send only the AID key.

Sendstring

Sendstring sets the string at the current cursor location and sends the data with the Enter key.

Parms String to be sent

Return

Returns 0 if successful, or -1 if there is an error or the terminal gets disconnected

Example:

```
Sendstring ("ctsoa"); // enters the string ctsoa at the
current cursor location and sends the data with the Enter key
```

Setcursor

Setcursor positions the cursor on the terminal.

Parms Screen position can be either of the following:

ROW,COL - the Row and Column to position the cursor Location - the relative location on the screen

Return

Returns the relative location of the cursor, or -1 if there is an error or the terminal gets disconnected

Example:

```
SetCursor(1,1); // sets the cursor to the top left position of the screen
SetCursor(80); // sets the cursor to offset 80 (This would be row 2,
col 1 on a 24x80 terminal)
```

SetDefaultTimeout

SetDefaultTimeout sets the default timeout.

Parms New wait time, in milliseconds. For example, 1000 is a one second wait.

Return

Returns the previous value of the WaitForScreen Timeout valued

Example:

```
oldTimeoutValue = setDefaultTimeout(10000);
// override the timeout to 10 seconds
```

SetField

SetField places a string in the field containing the designated location The number of characters is determined by the length of the string.

Parms String value - the string to be placed in the field

Position - a position on the screen that is part of a field.The position must be in an Unprotected field.This can be a ROW,COL or relative location on the screen.

Return

Returns 0 if successful, or -1 if there is an error.

Example:

```
SetField(applid,5,1); // put the string pointed to by applid
at the field containing row 5,col 1
SetField("ctsoa",80); // put the string ctsoa at the field
containing offset 80 (row 2 col 1 on a model 2)
```

Note: If the location on the screen is not an Unprotected field, or if the length exceeds the size of the field, an error will be generated; no data will be placed.

SetFontSize

Sets a new font size.

Parms The new font size, in points

Return

Returns the previous value of the font size

Example:

```
oldFontSize = setFontSize(16); // set the new Font size to 16 pt
```

Setstring

Setstring places a string at the designated location on the screen. The number of characters is determined by the length of the string.

Parms String value - the string to be placed on the screen:

Position - the position on the screen. This can be a ROW,COL or relative location on the screen.

Return

Returns 0 if successful, or -1 if there is an error

Example:

```
SetString(applid,5,1); // put the string pointed to
by the variable applid at row 5,col 1
SetString("ctsoa",80); // put the string ctsoa at
offset 80 (row 2 col 1 on a model 2)
```

Note: If the location on the screen is not an Unprotected field, or if the length exceeds the size of the field, an error will be generated and no data will be placed.

Show Shows the terminal screen. This sets the terminal frame to visible.

Parms None

Return

Returns the relative location of the cursor, or -1 if there is an error or the terminal gets disconnected

Example:

```
Show(); // set the terminal window to visible
```

WaitforScreen

WaitForScreen waits for a screen of data to be received from the host.

Parms arg 1 = Timeout value. This is the length of time to wait, in milliseconds.

If no data is received by this time, processing continues If no value is specified, a default of 15 minutes is assume

Return

Returns a numeric value specifying the result of the function:

0 = data was received

-1 = timeout. No data was received within the time specified

Example:

```
waitforscreen(); // wait until a screen is received
(or 15 minutes has elapsed)
waitforscreen(15*1000) // wait for up to 15 seconds
```

Note: Because the WaitforScreen is waiting for input, it could take a long
time before a screen is received. If a stall has occurred and the system is
not responding, the script will stall. For this reason, there is a timeout
value associated with this function. If no data is received for 10 minutes,
the wait is cancelled automatically and a value of -1 is returned.

WaitForString

WaitForString waits for a screen containing the specified data at the
specified location to be received.

Parms   arg 1 = The string to wait for

arg2 = The row,column position of the data

arg3 = Timeout value. This is the length of time to wait, in
milliseconds.

If no data is received by this time, processing will continue.

If no value is specified, a default of 15 minutes is assumed

Return
Returns a numeric value specifying the result of the function:

0 = data was received

-1 = timeout. No data was received within the time specified

Example:

```
Waitforstring("Userid",5,1); // wait until a screen is received
that has the string "Userid" at row 5 col 1
Waitforstring ("Userid",5,1 1000*15) // wait for up to 15 seconds
```

Note: Because the waitforstring is waiting for input, it could take a long
time before a screen is received. If a stall has occurred and the system is
not responding, the script will stall. For this reason, there is a timeout
value associated with this function. If no data is received within 10
minutes, the wait will be cancelled automatically and a value of -1 will be
returned.

Sample script

The following example shows a sample script for an Omegaview Zoom. It
demonstrates the constructs of the scripting language through user-defined
functions, local variables, flow control and built-in emulator functions.

```
/*
 Omegaview Zoom Emulator Script
  Copyright(c) !IBM Corporation 2002, 2005
*/
// -------------------------------------
// Check for Signon Panel
// 0 = not found
// 1 = found
// -------------------------------------
int FUNCTION signonPanel()
{
  if ((findString("OMEGAVIEW") > 0) AND
    (findString("Sign On Panel") > 0) AND
    (findString("Userid") > 0) AND
```

```
    (findString("Password") > 0) )
  return TRUE;
  else
  return FALSE;
}

// ---------------------------------------
// Check for Signon Error
//   return 0 if no error
//            1 if the SignonError pop-up found
// ---------------------------------------
int FUNCTION signonErrorpop-up()
{
  if ((findString("OMEGAVIEW") > 0) AND
      (findString("User access denied") > 0) )
  return TRUE;

  return FALSE;
}

// ---------------------------------------
// Goto Command Panel
// 0 = not found
// 1 = found
// ---------------------------------------
int FUNCTION gotoCommandPanel()
{
  rc = TRUE; // be optimistic

  if (findString(CMD_PANEL_ID) < 0)
  {
    sendkey(DEFHOTKEY_CODE);
    waitforScreen();
    if (findString(CMD_PANEL_ID) < 0)
        rc = FALSE; // not found

  }

  return rc;
}

void FUNCTION displayError(int errCode)
{
  if (errCode == 4)
    msgBox("Application not available");
  else if(errCode == 8)
    msgBox("Application not available - no signon Panel");
  else if (errCode == 16)
    msgBox("Invalid Signon request");
  else if (errCode == 24)
    msgBox ("Internal Error - command Panel not found");
  else
    msgBox ("Unknown Error. rc = " + errCode);
}
// ---------------------------------------
// LOGON
//   @param - user userid
//   @param - pswd password
//
//   @returns 0, success
//            4, no logo screen
//            8, no signon Panel
//           16, signon error
//           24, no command Panel


int FUNCTION logon(string user,string pswd)
```

```
{
rc = 0;
if (user == null || user == "")
 return 1;

sendkey("HOME");
sendkey("TAB");sendkey("TAB");sendkey("TAB");
sendstring("VGMVH"); // applid
waitforscreen();
waitforscreen(); // OV sends two screens (unlock)

while (1)
{
  if( findString("Press Enter to begin") < 0)
  {
    rc = 4;
    break;
  }
   // We have the Logo screen
  setCursor(1,1);
   sendkey("ENTER");
  waitforScreen();

  // Did we get the Signon Panel?
   if (signonPanel() == FALSE)
     {
      rc = 8; // no Logo
       break;
     }

   // do we have the F11 screen?
  if (findString("F11=CMW") < 0) // not found
  {
  // go the 4th field and do the magic
  sendKey("HOME");sendKey("TAB");sendKey("TAB");sendKey("TAB");
  setString("MVPM");
  sendKey("TAB");
  setString("MVPM");

   }
   else
  {
    sendKey("F11");
    waitforScreen();
    setstring("MVPM");
    sendKey("F24");
    WaitForScreen();
  }
  // now continue the logon
  sendKey("HOME");
  setString(user);
  sendKey("TAB");
  setString(pswd);
  sendKey("ENTER");
  waitforScreen();

  // if an error...
  if (signonErrorpop-up() == TRUE)
  {
    sendkey("F12"); // clear the pop-up
    waitForScreen();
    sendkey("F3"); // end the session
    waitForScreen();
    ;rc= 16; // signon error
    break;
  }
  rc = gotoCommandPanel();
```

```
      if (rc != 1)
        {
          rc = 24;
          break;
        }
    rc=0;
    break;
 }


 // success
 return rc;
}

// Main Program starts here

CMD_PANEL_ID = "KMVPMMAI";
DEFHOTKEY_CODE = "PA2";
EDITKEY_TAG = "=Cancel";
DEF_EDITKEY_TAG= "F12";

// prompt for the userid and password
// (can get this from system vars...)
userid = prompt("Enter UserID");
pswd = prompt("Enter Password","PSWD");

hide(); // don't show them the screen flicker

// perform the Logon
logon_complete = logon(userid,pswd);

if (logon_complete != 0)
{
  displayError(logon_complete);
  return;
}

// Logon was successful. So Zoom to the Omegamon
// (we're hard wire to IMS Workload
// until system parm passing is supported)
//
Setstring("KMVPMZM1");
sendkey("TAB");
setString("imsdc.wrt");
sendKey("ENTER");
waitforScreen();

show(); // show the results (should be the Omegamon screen)

// we don't want to show the Omegaview Command Screen
// so when they dismiss the Omegamon zoom, we will disconnect
while (1)
{
  waitforScreen(20000);
  if (findString(CMD_PANEL_ID) >= 0)
    {
      sendKey(DEF_EDITKEY_TAG);
      waitforScreen();
      break;
    }
 }
```

Chapter 7. Custom queries

In views that display monitored data, attribute values from managed systems are retrieved by queries to the Tivoli Enterprise Monitoring Server. You can edit the queries that are used in the predefined workspaces for your monitoring products, or create new queries to populate new views. In addition, you can retrieve data from any JDBC- or ODBC-compliant database to display in a chart or table by writing an SQL SELECT statement.

Understanding queries

Queries are made to the Tivoli Enterprise Monitoring Server to retrieve sampled attribute values from managed systems for display in workspace chart views, the table view and relational table-based topology views. There are benefits and limitations that you must be aware of when creating custom queries to the monitoring server or SQL queries to a JDBC- or ODBC-compliant data source.

Queries to the monitoring server

For queries written to the monitoring server, the Queries editor enables you to add a pre-filter, composed in the same way as a filter for a table or chart view in the Properties editor. Additionally, you can add a column function to average, count, total, or find the minimum or maximum value of a column grouping. The example below shows the Tivoli Enterprise Monitoring Server query `Open Events Count - Managed System`.

	fx **Local Timestamp**	*fx* **Status**	*fx* **Origin Node**	*fx* **Situation Name**	**Type**
1	☑	☐	☑	☑	☑
2		== Open	== $NODE$		
3	== 14,000				
4	== 14,999				

For most predefined queries, the specification shows == $NODE$ in the column that identifies the system (such as Server Name or Origin Node). The $NODE$ is a symbol that is replaced by the node name of the agent or agents associated with the Navigator item where the query is applied. This filter criteria is required to limit the retrieved data to only those agents associated with the Navigator item rather than all agents of that type.

You can clear the box below the column heading for any attribute you want to write a filter against, but which you do not want to retrieve. The example above has a filter written on the Status attribute to show only true situations (== Opened), although the attribute itself will not display in the view.

The query in the above example also shows the COUNT column function in the Type column. Although not shown in the query editor, the Situation Name was specified as the GROUP BY column (click **Advanced**). Tivoli Enterprise Portal will keep count of the number of times an event is opened and show the total in the Type column.

Here is an example of the queries for Managed System attributes. If you apply the Current Managed Systems Status query to it, all attributes are

retrieved; if you apply the Windows Systems Online or Windows Systems Offline query, only 4 attributes are retrieved. With either query you can use the Filters tab in the Properties editor to display fewer attributes.

🖳 Managed System	*fx* Managing System	*fx* Status	*fx* Origin Node	*fx* Name	*fx* Product	*fx* Version
🗊 Current Managed Systems Status	☑ == $AGENT$	☑	☑	☑	☑	☑
🗊 Windows Systems Online		☑ =='*ONLINE'		☑	☑ =='NT'	☑
🗊 Windows Systems Offline		☑ =='*OFFLINE'		☑	☑ =='NT'	☑

To determine which queries to use, consider that every attribute you retrieve adds to the processing time and network traffic, as does every query. If possible, use the same query for more than one view in a workspace and create queries that eliminate unwanted attributes or values.

Your IBM Tivoli Monitoring predefined workspaces use queries to the monitoring server to retrieve attributes from agents. You can also create your own queries with only the attributes you want to see in the view. The attributes can be from one group only; you cannot mix attributes from different groups in the same query.

Queries to JDBC and ODBC data sources
The Query editor has a text editor where you can compose free-form SQL queries to any JDBC- or ODBC-compliant databases located on the system where the Tivoli Enterprise Portal server is installed. This gives you the ability to integrate systems management data from monitoring agents with data from other sources, such as a third-party database, in one workspace.

Pre-filters compared with post-filters
Your monitoring product has queries for its attribute groups that send a request to the Tivoli Enterprise Monitoring Server for data from managed systems. With the Query editor, you can create your own queries to the monitoring server and add a *pre-filter* that is composed in the same way as a *post-filter* for a table, chart, or relational table based topology view. Additionally, you can add a column function to average, count, total, or find the minimum or maximum value of a column and group the results accordingly.

Advanced options
The Query editor has advanced options for sorting the view results, the attribute to group the rows by when a column function is used, and the number of rows to display from the beginning or end of the list of returned values (such as the top ten processes).

Queries to the monitoring server

Queries are made to the Tivoli Enterprise Monitoring Server to retrieve sampled attribute values from managed systems for display in workspace chart views, the table view and relational table-based topology views.

Creating a query to the monitoring server

You can create a new query to the monitoring server. Use this method to create a query only if you know what symbols the are required by the query.

Before you begin

Although you can follow this method to add a new query to the monitoring server, use 🗐 **Create Another Query** to ensure the query has the required symbol expression to reference the managed systems to include. Otherwise, no data is retrieved. **Exception**: When your product is upgraded to a new release with new attributes or attribute groups, the new attributes are not shown in the queries for that attribute group. And no new queries are shown for new attribute groups. If this exception applies, use this procedure to create a new query.

Your user ID must have permission for Modify Query and for Workspace Author Mode. If you do not see the 🗊 **Queries** tool, your user ID does not have View or Modify Query permissions. If you can see the tool but it is disabled, your user ID does not have Workspace Author Mode permission. If you can open the Query editor but the tools are disabled, your user ID does not have Modify Query permission.

Procedure

1. Complete one of the following steps to open the Query editor:
   - Click 🗊 **Queries**.
   - In the title bar of the table or chart view where you would like to apply a new query, click ✎ **Properties** > 🗊 **Click here to assign a query**.
2. Click 🗋 **Create New Query**.
3. In the **Name** field, type a name of up to 32 letters, numbers, underscores (_), and spaces; in the **Description** field, type a description to help users identify what the query does or what it is for, up to 256 characters and spaces. The Select attribute window opens.
4. Select the Tivoli product name from the Category ▾ list, and click **OK**.
5. Select the attribute group from the Group list, select the attributes to include from the Item list, and click **OK**. After selecting one attribute, you can use Ctrl+click to select multiple attributes or Shift+click to select all attributes from the first selection to this one, or click **Select All**.
6. Add the same symbol expressions to the specification as those used in the predefined query for this attribute group. Usually, this symbol is the $NODE$ symbol in the column that identifies the computer: System Name, Server Name, Originnode or something similar. If you are not sure what the expressions are, click **Apply**, open a predefined query for the attribute group that includes a symbol expression. Make note of the column name and any symbol expressions used in the specification, return to your new query, and add the symbol expression to the appropriate column.
7. Edit the specification for any changes you want to make:
   - Add an attribute to the query by clicking **Add Attributes** and selecting the attributes you want to include. The attributes available are from the group used in the original query.
   - Remove an attribute by right-clicking the column heading and clicking ✖ **Delete**. (You can also right-click a row and delete it.)
   - Include an attribute in the filter, but do not retrieve it by clearing ☐ that column.

- Clear the filter criteria by right-clicking the cell, column, or row, and clicking ◣ **Clear Contents**.
- Insert a row by right-clicking the row and clicking ⊡ **Insert**.
- Cut, copy, or paste by right-clicking the cell or row and clicking ✂ **Cut**, ▣ **Copy**, or ▣ **Paste**. When you paste, the contents of the clipboard overwrites the cell or row.
- Add filter criteria by clicking in a cell and writing the expression.
- Add a column function by clicking in the column heading and selecting a function from the list; then click **Advanced** and select the column to Group By.
- Specify a sort order by clicking **Advanced**, and selecting the column to Sort By. Note that if you selected a Group By column, you cannot also specify a sort order.
- Adjust column order by dragging a column heading and dropping at the insertion point.
- Restrict the rows (data series) that show in the table view or bar chart to a specific number. Restrict the first rows retrieved from the monitoring agent or the most recent (last rows) by clicking **Advanced**, selecting ◉ **First** or ◉ **Last**, and specifying the [10] ⊟ the number or rows to retrieve.

8. Optional. If you opened the Query editor from the view properties, you can change the list of managed systems the query will extract data from.

   a. Click the **Query Results Source** tab.
   b. Select ◉ **Let user assign explicitly**, then ▶ remove managed systems from the **Assigned** list and ◀ add them from the **Available** list.

   Regardless of the managed systems assigned to a query, a gauge chart can show data from only one managed system. If you are using the query in a workspace where multiple systems are applicable, such as at the ⊞ Enterprise- or ▤ Platform level of the Navigator Physical view, the data shown is from the first managed system in the **Assigned** list.

9. When you are finished, click **Apply** to save the query and keep the window open or click **OK** to save the query and close the window. If you opened the Query editor from the Properties window, clicking **OK** also selects the query for the current view.

What to do next

Be aware of these limitations:

1. The attributes in a query can be from one group only; you cannot mix attributes from different groups in the same query.

2. Even if data collection has been started, you cannot use the time span feature if the query for the chart or table includes any column functions. If this is the case, you can select or create another query to enable ▣ **Time Span**.

3. The **Sort by** advanced query function is incompatible with the historical reporting feature. The view will not be sorted even if you have specified a sort order in the query. Nonetheless, while a table view that specifies a time span is displayed in a workspace, you can still sort it by clicking a column heading.

   Similarly, the **Sort by** and **Group by** advanced query functions are unavailable for views that specify a time span within the last 24 hours, that is, views that draw from short-term history. (The first 24-hour retrieval from short-term history can be reduced or eliminated by adjusting the KFW_REPORT_TERM_BREAK_POINT value in the Tivoli Enterprise Portal Server environment file.)

Creating another query to the monitoring server

Use this method to create Tivoli Enterprise Monitoring Server queries and to create queries similar to existing ones.

Before you begin

Your user ID must have permission for Modify Query and for Workspace Author Mode. If you do not see the ▯ **Queries** tool, your user ID does not have View or Modify Query permissions. If you can see the tool but it is disabled, your user ID does not have Workspace Author Mode permission. If you can open the Query editor but the tools are disabled, your user ID does not have Modify Query permission.

If your product has been upgraded to a new release, with new attributes or attribute groups added, and you see no new attributes in the query or queries for that attribute group or you see no query for a new attribute group, do not use these steps. Instead, use the procedure for creating a new query.

Procedure

1. Complete one of the following steps to open the Query editor:
   - Click ▯ **Queries**.
   - In the title bar of the table or chart view where you would like to apply the new query, click ✎ **Properties** > ▯ **Click here to assign a query**.
2. Click ⊞ to expand the ◉ Monitoring agents and ▯ attribute group folders to see the queries.
3. Find a query similar to the one you want, select the query and click ▯ **Create Another Query**.
4. In the **Name** field, type a name of up to 32 letters, numbers, underscores (_) and spaces, then click **OK**. A duplicate of the original query is displayed with the name you entered.
5. In the **Description** field, type a new description.
6. Edit the specification for any changes you want to make:
   - Add an attribute to the query by clicking **Add Attributes** and selecting the attributes you want to include. The attributes available are from the group used in the original query.
   - Remove an attribute by right-clicking the column heading and clicking ✖ **Delete**. (You can also right-click a row and delete it.)
   - Include an attribute in the filter, but do not retrieve it by clearing □ that column.
   - Clear the filter criteria by right-clicking the cell, column, or row, and clicking ✎ **Clear Contents**.
   - Insert a row by right-clicking the row and clicking ▯ **Insert**.
   - Cut, copy, or paste by right-clicking the cell or row and clicking ✂ **Cut**, ▯ **Copy**, or ▯ **Paste**. When you paste, the contents of the clipboard overwrites the cell or row.
   - Add filter criteria by clicking in a cell and writing the expression.
   - Add a column function by clicking in the column heading and selecting a function from the list; then click **Advanced** and select the column to Group By.

- Specify a sort order by clicking **Advanced**, and selecting the column to Sort By. Note that if you selected a Group By column, you cannot also specify a sort order.
- Adjust column order by dragging a column heading and dropping at the insertion point.
- Restrict the rows (data series) that show in the table view or bar chart to a specific number. Restrict the first rows retrieved from the monitoring agent or the most recent (last rows) by clicking **Advanced**, selecting ○ **First** or ○ **Last**, and specifying the 10 ▪ the number or rows to retrieve.

7. Optional. If you opened the Query editor from the view properties, you can change the list of managed systems the query will extract data from.

   a. Click the **Query Results Source** tab.

   b. Select ○ **Let user assign explicitly**, then ▶ remove managed systems from the **Assigned** list and ◀ add them from the **Available** list.

   Regardless of the managed systems assigned to a query, a gauge chart can show data from only one managed system. If you are using the query in a workspace where multiple systems are applicable, such as at the ▦ Enterprise- or ▢ Platform level of the Navigator Physical view, the data shown is from the first managed system in the **Assigned** list.

8. When you are finished, click **Apply** to save the query and keep the window open or click **OK** to save the query and close the window. If you opened the Query editor from the Properties window, clicking **OK** also selects the query for the current view.

Related concepts:

Appendix A, "Formula functions," on page 415
Use the functions provided in these editors to write formulas: Situation, Query, Properties Filters tab, and Properties Thresholds tab.

Adding symbols in queries to the monitoring server

Most predefined queries include a == $NODE$ (or similar) variable expression that identifies the managed systems. The $NODE$ symbol resolves to the node name of each agent associated with the Navigator item and serves to limit the retrieved data to those agents rather than all agents of that type in your enterprise.

About this task

You can also add your own symbols (variables) to a query to provide more customization options in the Link Wizard. To check for or to add a symbol to a query, complete the following steps:

Procedure

1. Open the workspace to which you plan to link (the target, or link to, workspace).

2. Right-click the table or chart. This view has a key field you want to reference in the link definition. Typically, it is the view that will show details for an attribute from the row where you launched the link.

3. Click ▦ **Properties > 🗋 Click here to assign a query**.

4. In the Query editor, review the query specification to see that it has a symbol defined for the attribute on which you want to build the custom link. Symbols are surrounded by dollar signs ($) .

5. If the attribute has no symbol, create your own by typing $mysymbol$ in the first cell in that column, where mysymbol is a one-word name. If the query

already has a symbol (or symbols) defined, such as `$WTPROCESS.IDPROCESS$`, be sure to add your symbol to the same row of the specification to create an AND expression.

You can create a symbol for any column whose value can be used to restrict the number of rows returned. If the symbol resolves to a non-null value, a restricting where clause is added; if not, it is skipped.

6. Click **OK** to save your changes to the query and close the window.
7. Repeat these steps for any other views whose query you want to check.

Results

When you are finished, you can open the source workspace for the link and start the Link Wizard from the chart or table row. The symbol you added is displayed in the list of symbols available for the query.

- Use a one-word name surrounded by dollar signs, as in $mysymbol$, where mysymbol is the symbol name.
- Do not use the same name as one of your monitoring agent attributes. For example, $MYNODE$ is acceptable, but not $NODE$.
- When the link is launched, a symbol is resolved to a value enclosed in single quote marks if it is anything other than an integer. An integer resolves without the quotes.

Queries to a database

The Query editor has a text editor where you can compose free-form SQL queries to any JDBC- or ODBC-compliant databases located on the system where the Tivoli Enterprise Portal server is installed.

Creating a query to a JDBC or ODBC database

The Query editor has a text editor for composing free-form SQL statements to any JDBC- or ODBC-compliant database located on the computer where the Tivoli Enterprise Portal server is installed.

Before you begin

Your user ID must have Modify Query permission for this function. If you do not see the Queries tool, your user ID does not have View or Modify Query permissions; if you can open the Query editor but the tools are disabled, your user ID does not have Modify Query permission.

About this task

Use this task to create a new query of JDBC or ODBC attributes, thus enabling you to integrate systems management data from monitoring agents with data from other sources in one workspace.

Procedure

1. Do either of the following actions to open the Query editor:
   - Click ⬚ **Query editor** in the main toolbar.
   - In the title bar of the view where you want to apply the new query, click ✎ **Properties >** ⬚ **Click here to assign a query**.
2. Click ⬚ **Create New Query**.

3. In the **Name** field, type a name of up to 32 letters, numbers, underscores (_) and spaces and, optionally, in the **Description** field, type a description of up to 256 characters and spaces.

4. From the **Category** ▾ list, select the folder where you want the new query to appear in the Queries tree.

5. Select a JDBC or ODBC database in the **Data Sources** list. ☑ Custom SQL is selected for you. If you do not see the database for which you want to create the query, it must be added through Manage Tivoli Monitoring Services.

6. Click **OK**. The new query is displayed in the Queries tree under ⊚ *Folder* (where *Folder* is the Category you selected) in the 🖳 Custom_SQL folder. The **Specification** tab opens with a Custom SQL text box for you to enter SQL commands.

7. Write an SQL SELECT statement in the text box using the correct syntax. You must be familiar with SQL commands and syntax to write a valid SQL query.

8. When you are finished writing the SQL statement, click **Apply** to save the query and keep the window open; or click **OK** to save the query and close the window.

What to do next

If the view remains empty after you assign the query, the query has not been entered correctly. Review the notes for JDBC or ODBC database queries and return to the Query editor to edit the SELECT statement as needed.

Notes on queries to a JDBC or ODBC database

Consult these notes about SQL statements written to a JDBC or ODBC database to ensure data is returned successfully.

Date and time format varies among databases
You must know the format and write the appropriate SQL code against it when selecting by date. Otherwise, the statement might be not be valid for the database to which the request is directed.

Database user IDs
Anyone with permission to create custom queries has access to all the JDBC and ODBC DSNs (Data Source Names) that were created on the same computer as the Tivoli Enterprise Portal Server. Add database user IDs (to be used in the DSN) to your database software to restrict user access to only those tables, columns, and so on, allowed by your organization's security policies.

Symbolic variables
If you included a symbol in the statement, you are prompted to enter a value for the symbol when you apply the query to a table or chart view. The value you enter filters the data that is displayed in the view, but is temporary and cannot be saved with the workspace.

Column functions in a statement that has a JOIN
When a custom SQL statement contains a JOIN command, the Tivoli Enterprise Portal does not set the table name property in the request sent to the Tivoli Enterprise Portal Server. Typically, the portal server determines the table name from the result set returned from the database. However, when a column function (such as COUNT, MAX, or MIN) is used on one of the columns in the request, the result set meta data does

not provide a table name for that column. With no table name, the portal client treats the column as a text attribute, which means that it cannot be rendered in a chart view.

To ensure that the column data type is set correctly, provide an alias for any column in the SELECT statement that uses a column function. The alias must be in the form **table.column**. The table and column values can be arbitrary strings, and do not need to match the actual table or column name in the database. The value supplied for **column** is the column heading. In the attribute selection dialog, the value for **table** is the attribute group, and the value for **column** is the attribute item.

In the **Before** sample, the data returned from the statement renders properly in a table view but cannot be selected in the ▦ Properties editor ↷ Filters tab for any chart views because the returned values are read as text attributes. The **After** sample shows how the statement was edited to create an alias to explicitly state the table name for TYPE and COUNT.

Before
```
SELECT COUNT(column1) item_count
FROM table1
LEFT OUTER JOIN table2 ON table1.column1 = table2.column1
```

After
```
SELECT COUNT(column1) "my_group.item_count"
FROM table1
LEFT OUTER JOIN table2 ON table1.column1 = table2.column1
```

Using * in a statement that has nested queries

When a custom SQL statement contains * in the parent query of the nested query, the Tivoli Enterprise Portal does not set the table name properly in the returned result set. Typically, the portal server determines the table name from the result set returned from the database. However, when * is used on the parent query in the request, the result set meta data does not provide a table name for that column.

Instead of using * in the parent query, provide an alias for any column in the SELECT statement that uses a column function. The alias must be in the form **table.column**. The value supplied for **column** is the column heading. In the attribute selection dialog, the value for **table** is the attribute group, and the value for **column** is the attribute item.

In the **Before** sample, the data returned from the statement renders properly in a table view but cannot be selected in the ▦ Properties editor ↷ Filters or Thresholds for any table views. The **After** sample shows how the statement was edited to create an alias to explicitly state the table name.

Before
```
select * from ( select column1, coulumn2 from table1 )
```

After
```
select column1 "table1.column1" , column2 "table1.column2"
from ( select column1, coulumn2 from table1 )
```

Creating another query to a JDBC or ODBC database

If you want to create a query that is similar to an existing one, you can copy the original and edit the new query.

About this task

To create another query to an ODBC database:

Procedure

1. Do one of the following to open Query editor:
   - Click 🔲 **Queries**.
   - In the title bar of the table or chart view where you would like to apply the new query, click ✏ **Properties** > 🔲 *Click here to assign a query*.
2. Click ⊞ to expand the 🔵 Tivoli Enterprise Monitoring Agents and 🖥 attribute group folders to see the queries.
3. Find a query similar to the one you want, select the query and click 🔲 **Create Another Query**.
4. In the **Name** field, type a name of up to 32 letters, numbers, underscores (_) and spaces, then click **OK**. A duplicate of the original query is displayed with the name you entered.
5. In the **Description** field, type a new description.
6. Edit the specification for any changes you want to make to the SELECT statement.
7. When you are done, click **Apply** to save the query and keep the window open; -OR- click **OK** to save the query and close the window.

Symbols in creating queries to JDBC or ODBC data sources

You can use symbols (variables) in custom SQL queries to JDBC- or ODBC-compliant data sources. The reason you would add a symbol to a custom SQL query is to build a filtered link (using the Link Wizard) to the workspace with a table or chart that uses this query. For example, if you are building a link based on the data source attribute for system name, the symbol might be $myhostname$.

When you apply the query to a table or chart, you will be prompted to enter a value for the variable as soon as you click **Apply** or **OK**. The value filters the data that will show in the view, but is temporary and cannot be saved with the workspace.

- Use a one-word name surrounded by dollar signs, as in $mysymbol$, where mysymbol is the symbol name.
- When the link is launched, a symbol is resolved to a value enclosed in single quotation marks if it is anything other than an integer. An integer resolves without the quotation marks. If you are referencing the symbol in a custom SQL statement and you do not want the quote marks included in the resolved value, add a pound sign after the first dollar sign, as in $#noquotes$, where noquotes is the symbol name. This notation is required for numeric data types.

  Example:

  ```
  SELECT * FROM MYTABLE WHERE MYCOLUMN = $#myNumber$
  ```

  The value of the $#myNumber$ symbol is a number (such as 5 or 345).

Adding a database to the Data Source list

The Data Source area of the Create Query window lists the databases available for a query. The databases that are shown are read from the database server on the system where the Tivoli Enterprise Portal Server was installed. If you do not see the database for which you want to write a query in the list, you must add it through Manage Tivoli Monitoring Services.

About this task

You can also use the **tacmd** **configurePortalServer** to complete this task. See the *IBM Tivoli Monitoring Command Reference* for more information.

Restriction: The ODBC interface is not supported on Linux and UNIX. Because of this limitation, the `CONNECTION_TYPE=DB2` is used.

To add a database to the Data Source list on Windows:

Procedure

1. On the computer where the Tivoli Enterprise Portal Server is installed, click **Start > Programs > IBM Tivoli Monitoring >** **Manage Tivoli Monitoring Services**.
2. Right-click the Tivoli Enterprise Portal Server name and click **Advanced > Edit ENV File**.
3. Find the section entitled KFW_ODBC_*.
4. Delete the # remark symbol at the beginning of the DSUSER1 entry.
5. Edit the DSUSER1 entry for your data set name, username and password:

   `DSUSER1="DSN=myODBC;UID=mUsername;PWD=myPassword"`

   where
   - `DSN` is the ODBC data source name as it is displayed in the database server.
   - `UID` is the user ID for accessing the database.
   - `PWD`, if used, is the password for the user ID.

   Each property is a key=value pair and is separated by a semi-colon (;). Other key=value pairs can be added such as: `CONNECTION_TYPE`, `CONNECTION_LIMIT`, `KFWJDBCDRIVER`, and `KFWDSURL`. The information provided for the entry creates your datasource connection string.
6. Continue to add new lines for other data sources, from DSUSER2 up to DSUSER9.
7. Close the text editor and answer **Yes** to the prompt to recycle the portal server.

   Note: The database you want to reference must be in the ODBC Data Source registry on the Tivoli Enterprise Portal Server. If it is not, consult your Windows system administrator.
8. Recycle the portal server.

Editing a query

You can edit any of the queries you have created for the monitoring server or for an ODBC-compliant database. You might want to modify a query to add a filter that keeps unnecessary data from cluttering the view and slowing the return of data, for example.

Before you begin

Your user ID must have permission for Modify Query and for Workspace Author Mode. If you do not see the **Queries** tool, your user ID does not have View or Modify Query permissions. If you can see the tool but it is disabled, your user ID does not have Workspace Author Mode permission. If you can open the Query

editor but the tools are disabled, your user ID does not have Modify Query permission.

Procedure

1. Complete one of the following steps to open the Query editor:
   - Click **Query Editor**.
   - In the view whose query you want to edit, click **Edit Properties**.
   - Click **Properties** and select the view from the Properties tree.
2. Expand the Tivoli Enterprise Monitoring Agents and attribute group folders to see the queries.
3. Select the Query that you want to edit. The query specification is displayed in the right frame.
4. Edit the specification for any changes you want to make.
   - For monitoring server queries:

     Add an attribute to the query
     > Click **Add Attributes** and select the attributes to include. The attributes available are from the group used in the original query.

     Remove an attribute
     > Right-click the column heading and click **Delete**. You can also right-click a row and delete it.

     Include an attribute in the filter but not retrieve the value
     > Clear the check box for that column.

     Delete the filter criteria
     > Right-click the cell, column, or row and click **Clear Contents**.

     Insert a row
     > Right-click a row and click **Insert**.

     Cut, copy, or paste
     > Right-click the cell or row and click **Cut**, **Copy**, or **Paste**. When you paste, the contents of the clipboard overwrites the cell or row.

     Add filter criteria
     > Click in a cell and write the expression.

     Reorder the columns
     > Drag a column heading and drop it at the insertion point. Views that use this query are displayed with this column order. Users can drag columns in the workspace view to rearrange them and, if the workspace is saved the new order becomes permanent until another query is applied to the view or the columns reordered and the workspace is saved again.

     Specify a sort order
     > Click **Advanced** and select the **Sort By** column from the list. Note that if you selected a **Group By** column, you cannot also specify a sort order.
     >
     > If your environment has multiple monitoring servers, the sorting is done independently at each monitoring server and then returned to the portal without resorting the merged values. You can sort the rows returned after they are displayed in the view.

Specify the first or last *n* rows to return

Click **Advanced**, select ◎ **First** or ◎ **Last**, then use the [10 ⬍] spin control to adjust the number of rows to retrieve or type the number directly into the text box.

This is the number of rows to retrieve from *each* Tivoli Enterprise Monitoring Server in your monitored environment. For example, if your environment has a hub monitoring server with one remote monitoring server and you select the last 5 rows, the view shows 10 rows of data: 5 from the hub and 5 from the remote. The data is sorted for each server individually. In table views, you can click a column heading to sort the results from all servers.

Apply a group function

Click *fx* Formula in the column heading and select from the list of available group functions. The icon in the column heading changes to the icon for the chosen function:

▯ Minimum in group

▤ Maximum in group

▦ Count of group members

Σ Sum of group

▤ Average of group

Click **Advanced** and select the **Group By** column from the ⯆ list.

- For ODBC queries, edit the SELECT statement.

5. Optional. If you opened the Query editor from the view properties, you can change the list of managed systems the query will extract data from.

   a. Click the **Query Results Source** tab.

   b. Select ◎ **Let user assign explicitly**, then ▶ remove managed systems from the **Assigned** list and ◀ add them from the **Available** list.

   Regardless of the managed systems assigned to a query, a gauge chart can show data from only one managed system. If you are using the query in a workspace where multiple systems are applicable, such as at the ▦ Enterprise- or ▭ Platform level of the Navigator Physical view, the data shown is from the first managed system in the **Assigned** list.

6. When you are finished, click **Apply** to save the query and keep the window open or click **OK** to save the query and close the window. If you opened the Query editor from the Properties window, clicking **OK** also selects the query for the current view.

Related concepts:

Appendix A, "Formula functions," on page 415
Use the functions provided in these editors to write formulas: Situation, Query, Properties Filters tab, and Properties Thresholds tab.

Deleting a query

When you no longer need a query for any purpose, you can clean up the Query editor tree by deleting it.

About this task

▣ If you do not see the Queries tool, your user ID does not have View or Modify Query permissions. If you can open the Query editor but the tools are disabled,

your user ID does not have Modify Query permission.

Procedure

1. Do one of the following to open the Query editor:
   - Click 🔲 **Query Editor**.
   - In the view whose query you want to delete, click ✎ **Properties**.
   - Click 🖼 **Properties** and select the view from the Properties tree.
2. Select the 🔲 Query name in the tree. The query specification is displayed in the right frame.
3. Click ✖ **Delete Query**. If the Delete Query tool is disabled, your user ID does not have Modify Query permission.

   If a Navigator item is associated with a query, you will not be able to delete the query.

   Note: If you delete a query that is being used by a view, the view is no longer associated with a query and contains no data when you next open its workspace.
4. Click **OK** to confirm.

Tivoli Enterprise Monitoring Server predefined queries

Every IBM Tivoli Monitoring product is shipped with a set of predefined queries that you can apply to a table or chart view. This topic describes the queries delivered as part of the Tivoli Management Services infrastructure. The Tivoli Enterprise Monitoring Server queries are available for every item in the Navigator.

Expand the ⬤ **Tivoli Enterprise Monitoring Server** branch of the Queries tree to see these query groups:

🖥 **Agent Operations Log**
> This log and the EIB Change Log are good diagnostic tools to use when problems arise such as a true situation that is not raising an alert as expected. They also provide information that can help in planning optimization strategies.
>
> 🔲 *Agent Operations Log* provides information about situation and connectivity messages that occur on the managed system associated with the Navigator physical view item or the link source. It is useful for diagnostic research with IBM Support.

🖥 **Catalog Check**
> The Catalog Check queries compare the application catalogs of the hub Tivoli Enterprise Monitoring Server with those of the remote monitoring servers and report any discrepancies. These queries are used in the Manage Tivoli Enterprise Monitoring Servers workspaces:
>
> 🔲 *Compare Enterprise Catalogs*
>
> 🔲 *Compare TEMS Catalogs*

🖥 **Custom SQL**
> The custom SQL (Structured Query Language) queries provide information that the portal server (TEPS) database gathers from the hub Tivoli Enterprise Monitoring Server. These queries are for the Situation Event Results workspace and the Self-Monitoring Topology workspace.

EIB Change Log

The Enterprise Information Base log, like the Agent Operations log, is stored on the hub monitoring server.

EIB Change Log provides information about changes to Enterprise Information Base objects such as situations. It is useful mainly for diagnostic research with IBM Support. This is an ODBC query of the EIB Change Log attributes.

ITM Audit Log

The ITM Audit Log query is stored on the hub monitoring server.

ITM Audit Log provides information on critical state-changing events in the monitoring environment.

Managed System

The Managed Systems queries are used to check vital heartbeat information from the monitoring server and monitoring agents.

Current Managed Systems Status, when applied to a table view, displays all the Managed Systems attributes. It displays such information as the agent product code and the host name and IP address of the managed system. The Enterprise Navigator item has a Managed System Status workspace that uses this query. A threshold was applied to the Status column that highlights any offline managed systems.

ITV Managed System Count is a count of the managed systems reporting to each monitoring server.

Tivoli Enterprise Monitoring Servers

Windows Systems Online lists the Windows OS managed systems that are online. Other OS agents have similar queries.

Windows Systems Offline lists the Windows OS managed systems that are unavailable. Other OS agents have similar queries.

Situation Definition

Situation Definition uses the Situation Definition attributes to retrieve information that describe every aspect of the situations defined to this hub monitoring server.

Situation State Events

The Situation State Events queries collect status information on situation events.

Events by Owner Using ID

Selected Event Using ID

Similar Events by Situation Name lists all the events for the same situation. Such events would occur when the situation is true on more than one managed system or the situation is true on the same managed system more than once (a display item was defined in the Advanced Situation options window). You can see table views that use this query: From the Enterprise Status workspace, right-click the ✎ link anchor for an acknowledgement in the My Acknowledged Events view, and click Event Details - Similar by Situation Name.

Similar Events by Source lists all the events from the same managed system as the event you are viewing. You can see table views that use this

query: From the Enterprise Status workspace, right-click the ✐ link anchor for an acknowledgement in the My Acknowledged Events view, and click Event Details - Similar by Source.

📖 *Similar Events by Resource* lists all the events for situations that were defined with the same attribute group as the one used for this situation. You can see table views that use this query: From the Enterprise Status workspace, right-click the ✐ link anchor for an acknowledgement in the My Acknowledged Events view, and click Event Details - Similar by Resource.

📖 *Events by Owner Using Name* returns all the events by owner.

📖 *Selected Event Using Name* returns all events by their short name.

📖 *Similar Events by Situation Display Name* return all events by their full name.

📖 *Similar Events by Source using Situation Display Name*

📖 *Similar Events by Resource using Situation Display Name*

🖥 **Situation Status**

📖 *Situation Status Detail*

📖 *Situation Status Summary*

🖥 **Situation Status Current**
These queries use the Active Situation Status attributes to display status information on situations and policies that are running.

📖 *Alternate Current Situation Status* is an alternate query for message log initialization.

📖 *Current Active Situation Status Using ID*

📖 *Pure Event Occurrences Using ID* displays the situations with pure events opened when applied to a table view.

📖 *Current Active Situation Status Using Name*

📖 *Pure Event Occurrences Using Name* displays the situations with pure events opened when applied to a table view.

🖥 **Situation Status Log**
These queries use the Situation Status attributes to display status information on started situations and policies.

📖 *Open Events Count Using ID - Enterprise* shows the every open event on the Enterprise and the number of intervals the situation has been true since it was started.

📖 *Open Events Count Using ID - Managed System* shows the every open event on the managed system and the number of intervals the situation has been true since it was started.

📖 *Situation Status Log* shows the status of the situations on this managed system. It is similar to the message log view.

📖 *Open Events Count Using Name - Enterprise* shows the every open event on the Enterprise and the number of intervals the situation has been true since it was started.

Open Events Count Using Name - Managed System shows the every open event on the managed system and the number of intervals the situation has been true since it was started.

Chapter 8. Historical reporting

Tivoli Enterprise Portal provides historical data collection and reporting capabilities. If historical data collection has been configured and started for an attribute group on the managed systems specified, you can set a time span for retrieving historical data into a query-based view.

Overview of historical data collection and reporting

Use historical data collection and reporting to gather useful metrics about your managed network. You can also use historical data with the chart baselining tools for predictive analysis and in situation modeling for key performance indicators.

How historical data collection and reporting works

After historical data collection has been configured and started, data samples are collected and saved in history files at either the monitoring agent or the monitoring server for the short term. The data is stored in tables, one for each attribute group for which data is being collected. If the Tivoli data warehousing facility has been set up, you can move the data to a relational database for longer term storage.

The History Collection Configuration window enables you to configure historical data collection for individual attribute groups and specify from which managed systems to collect, how frequently to collect the data, where to store it for the short-term, and whether to save data long-term to a data warehouse. You can create multiple historical collections for an attribute group, each with different collection settings and managed system distributions, and you can write filtering formulas to collect only the data samples that meet the filter criteria.

The Object Group editor enables you to combine historical data collections into groups and assign managed systems. With historical collection groups you can easily turn on and turn off historical data collection at specific managed systems or groups of managed systems by changing the distribution.

Viewing historical data

After data collection has been configured and started for an attribute group, historical data for that group can be retrieved and displayed in query-based views (chart views, table view, and the topology view).

You know that historical data collection has been enabled for the attributes in a view when you can see ▣ **Time Span** in the view's toolbar. The view shows current data samplings unless you use this tool to specify a broader time period. The first 24 hours of collected data is retrieved from the short-term history files; beyond 24 hours, the data is pulled from the data warehouse.

Historical data is valuable for use in chart baselining and situation modeling. Applying historical data to chart baselines and situation modeling enables you to perform predictive analysis and fine tune situations.

Requests that return large amounts of data can negatively affect the performance of monitoring servers, monitoring agents, and your network.

For attributes or views that return considerable amounts of data, try to
limit the time span to as short a period as will provide the required
information.

Be aware that if this is a newly configured historical data collection, you
will see the historical data samples that have been collected and no more.
For example, if you created a historical collection and assigned it to a
historical collection group an hour ago, you will see only the last 1 hour in
the chart or table view even if you selected, say, the last 8 hours.

Historical navigation mode

The ▣ **Set historical navigation mode** tool at the end of the main toolbar is
for a different purpose: It enables you to set a time span that is
synchronized so that you can see the same time range in workspaces as
you move among them, and compare those values.

Related tasks:

"Setting a time span to display" on page 44
You can have the Tivoli Enterprise Portal log data samplings into history files or a
database for display in a table or chart. These historical data enabled views have a
tool for setting a range of previous data samples to be reported.

Setting a time span to display

You can have the Tivoli Enterprise Portal log data samplings into history files or a
database for display in a table or chart. These historical data enabled views have a
tool for setting a range of previous data samples to be reported.

Before you begin

Historical data collection must be configured and distributed to the managed
systems that you are querying data from. Otherwise, ▣ **Time Span** is not displayed
in the toolbar of the query-based view.

Some attributes groups, such as Situation Status and the Windows Event Log, are
historical in nature and show all their entries without you having to specify a time
span. For these types of attribute groups, you do not need to configure history
collection unless you want to roll off to a data warehouse for long-term storage or
limit the amount of data displayed. Thus, the time span feature, rather than
showing more data, limits the data reported to the time period indicated.

Even if data collection has been started, you cannot use the time span feature if the
query for the chart or table includes any column functions. If this is the case, you
can select or create another query to enable ▣ **Time Span**.

About this task

Take these steps to broaden the time range of data beyond the current data
samplings.

Procedure

1. Open the workspace containing the chart, table, or relational table-based
   topology view where you want to see historical data.
2. Click ▣ **Time Span** in the view's toolbar.
3. Select a time frame: **Real time plus Last _ Hours** (enabled for bar, plot, and
   area charts only), **Last _ Hours** (or Days, Weeks, or Months, if the data is
   warehoused), or **Custom**.

4. If you selected **Last** or **Custom**, specify the range of data:

- Detailed data is all the data collected for the agent.
- Summarized data is data that is aggregated across the specified time frame. Set the time frame interval, work shift, and the days. This option is disabled if no data warehouse and summarization schedule were configured for this attribute group.
- To set a custom time frame, click inside the **Start Time** and **End Time** fields to open the date editor. Use the ⬍ spin boxes to adjust the time, year, or month; and click the calendar day. **HH:MM:SS AM/PM** is initially set to the current time.
- To apply the time span to all views that use the same query as this view, select ☐ **Apply to all views associated with this view's query**. When this option is enabled, the query is modified to include the time span set here, so any other views using this query report the same time range.
- The **Timestamp** column that is added to the historical view can show either the global timestamp (the default) or local. Select ☐ **Use Hub time** to reflect the time at the Hub Tivoli Enterprise Monitoring Server rather than at the Tivoli Enterprise Portal Server or portal client.

Results

After you click **OK**, the view shows data from the time span you specified. If the view is a table, a timestamp is displayed as the first column and is accurate to the nearest minute; seconds are displayed as 00.

If you see `null` as the value of a table cell or chart point, it means that no value was stored in the Tivoli Data Warehouse. This happens when values that were identified as invalid are reported from a monitoring agent for a given summarization period.

What to do next

The sort function is incompatible with the historical reporting feature. If you are using ⬚ **Time Span** to retrieve historical data, the chart or table is not sorted even if you have specified a sort order in the query. You can still sort a table by clicking a column heading.

After support for an updated product has been applied to the portal server, it is possible to get a request error message about a missing or unknown column name in the view's status bar after you set a time span with ☑ **Use summarized data** selected. Wait until after the next scheduled summarization and pruning procedure has taken place before viewing the summarized data. If need be, you can reschedule summarization and pruning to run sooner.

Related concepts:

"Overview of historical data collection and reporting" on page 217
Use historical data collection and reporting to gather useful metrics about your managed network. You can also use historical data with the chart baselining tools for predictive analysis and in situation modeling for key performance indicators.

Related tasks:

"Discontinuing historical reporting"
For views that have a time span set to show data samplings over a range of time, you can remove the time span and display real-time data only.

Discontinuing historical reporting

For views that have a time span set to show data samplings over a range of time, you can remove the time span and display real-time data only.

About this task

To change the view to show current data samplings, perform the following steps:

Procedure

1. Open the workspace with the view from which you want to remove historical reporting.
2. Click **Time Span** in the view's toolbar.
3. Select ◉ **Real time**.
4. Click **OK**.

Results

If this is a table view, the leftmost time stamp column is no longer displayed. The view now shows data from the current sampling only. The exception is views of attribute groups that are normally historical, such as the Situation Status and the Event Log: When **Real time** is selected, attribute groups show all data they have collected, up to 2000 of the most recent entries.

Related tasks:

"Setting a time span to display" on page 44
You can have the Tivoli Enterprise Portal log data samplings into history files or a database for display in a table or chart. These historical data enabled views have a tool for setting a range of previous data samples to be reported.

Turning on historical navigation mode

Use historical navigation mode to synchronize workspaces and retain the same time period as you switch from one workspace to another.

About this task

Whenever you open a workspace, the latest values are retrieved to the query-based views. Open or link to another workspace and the data is retrieved at a different time. For performance analysis, you can specify a time span to use for the workspaces that get opened. Historical navigation mode is easily started and stopped, and you can adjust the time span. Follow these steps to start historical navigation mode.

Procedure

1. Do either of the following:
    *   Click ⬚ **Set historical navigation mode** at the end of the main toolbar.
    *   Click ⬚ **Select time span** in the view toolbar. The time span tool shows for views of attribute groups that historical data collection has been configured for and started or that are historical in nature.
2. Set the time span.
3. Select ☑ **Lock time span for Historical Navigation**.
4. If you want to base the time span on the Hub monitoring server's clock and time zone, select ☑ **Use Hub time**.
5. Click **OK**. From now on, workspaces that are opened will display data according to the time span settings, as indicated by the appearance of the Reset historical navigation mode tool. The status bar also shows an indicator of the time span.
6. When you are ready to deactivate historical navigation mode, click ⬚ **Reset historical navigation mode**, then clear ☐ **Lock time span for Historical Navigation**.

What to do next

Not all workspace views are able to display a time span or the range that you have specified. In this case, no values are displayed and a message in the view's status bar says that the query assigned to the view does not support a historical time range. It is likely that historical data collection has not been configured for the attribute group shown in the view or there is no historical data to report. Exit historical navigation mode to remove this message. To see historical data in this view, you must first configure historical data collection for the attribute group referenced by the query, then start collection and wait a sufficient period of time for historical data to be gathered before displaying the view again.

Historical collection configuration

To make historical data available for reporting in workspace views, you must configure and start historical data collection for the appropriate products and attribute groups. This is done through the History Collection Configuration window.

Understanding historical configuration

It is necessary to configure historical data collection for attribute groups and distribute to managed systems before users can see historical data in the Tivoli Enterprise Portal. Read this topic to understand the configuration settings and their implications.

Configure historical data collection first
> Customization of the Tivoli Enterprise Portal is best begun with configuring historical data collection. Until then, no historical data is available for the predefined historical workspaces that your product might offer, for situation modeling, or for chart baselining, all of which enable you to perform in-depth analysis and health assessments of your managed environment.

Distribute to Managed System (Agent) or Managing System (TEMS)
> In a historical collection, the first decision made regarding distribution is the method:

Managed System (Agent) is the default method and requires that any managed system in the distribution connect to a Tivoli Enterprise Monitoring Server Version 6.2.2 or later. The distribution goes to a subset of managed systems: the managed systems of that agent type that are assigned individually or as part of a managed system group. Alternatively, you can choose to assign managed systems for distribution to a historical configuration group that the collection belongs to.

Historical configuration groups provide a way to group historical collections that have the same distribution requirements. The managed system distribution list can be assigned for the group rather than to each collection individually.

Managing System (TEMS) is the method that was used for distribution in IBM Tivoli Monitoring Version 6.2.1 and earlier; it is the required method for distribution if the managed system connects to a V6.2.1 or earlier monitoring server. The distribution is to managed systems of that agent type that are connected to the Tivoli Enterprise Monitoring Server. If the Managing System (TEMS) method has been chosen for a collection definition, that collection becomes ineligible for membership in a historical configuration group.

Upgraded historical configuration from an earlier version
If you have upgraded to Tivoli Management Services Version 6.2.2 or later from a release prior to Version 6.2.2, you get a historical collection definition for each attribute group that was configured and the distribution method is unchanged: **Managing System (TEMS)**. If you would like to use the distribution techniques that are available when distribution is by managed system, change the distribution method for each collection definition to **Managed System (Agent)**. (See the link to "Changing the distribution method" at the end of this topic.)

Managed System distribution through the historical collection or through the historical configuration group
You can assign a managed system (or managed system group or both) to the individual historical collection definition, a historical configuration group or to a managed system group that is assigned to the collection or historical group.

As a best practice, develop a way to organize managed system distribution for historical data collection. It is possible for a managed system to be referenced more than once for the same historical collection definition. For example, the managed system might be have been assigned in the Distribution tab for the collection definition and the collection is a member of a historical configuration group that also has that managed system assigned. You need not be concerned with data being collected multiple times on that managed system for this collection definition, but be aware that if you intend to stop data collection on that managed system by removing it (or a managed system group it belongs to) from the distribution, data collection will continue to occur. This will go on until the managed system is removed from all group distributions that the collection is a member of.

Distributing a historical collection to a new managed system
After a new managed system is added to your managed environment, data collection begins immediately if the new managed system is assigned to a managed system group that a collection definition is distributed to or if the

data collection method is to **Managing System (TEMS)** and the new managed system connects to the Tivoli Enterprise Monitoring Server that was specified.

For example, a new Linux managed system is automatically assigned to the *LZ_SYSTEM predefined managed system group. If the historical collection is distributed to *LZ_SYSTEM, then historical data collection begins on the new managed system. Otherwise, you must add the new managed system to the historical collection distribution directly, to a managed system group that the collection is distributed to, or to the distribution of a historical configuration group that the historical collection is a member of.

Effect of editing a managed system group on the historical collection

When a managed system group gets updated, which is typically when a constituent managed system is removed or a managed system is added, any historical collection that includes that group is redistributed to all managed systems in the group and collection is restarted.

The redistribution and restarting of data collection can result in some gaps in the data if the collection interval is frequent, because collection is aligned with time boundaries. For example, if the collection interval is 15 minutes and collection was started at 3:16, the first sampling is at 3:30; or if the last collection was at 4:15 and an update to the managed system group causes the 4:30 collection to be missed, the next data collection takes place at 4:45.

Multiple collection definitions for the same attribute group

You can create another copy of a collection definition for an attribute group, then configure the copy for different values for any of these criteria: collection interval, warehouse interval, managed system distribution, or attribute filtering. What remains the same for every historical collection that is defined for an attribute group, however, is the Collection Location (TEMA or TEMS) and the settings for Summarization and Pruning.

To illustrate some of the possibilities, here is an example of three collection definitions for the Process attribute group:

- **Process transaction servers** has a 1-minute collection interval, a 1-hour warehousing interval, and is distributed to a few critical servers.
- **Process database servers** has a 15-minute collection interval, data warehousing every 12 hours, and is distributed to database servers.
- **Process staff systems** has a 1-hour collection interval, data warehousing every 24 hours, and is distributed to a general use group of managed systems.

The original and the two copied collection definitions have their short-term history files saved at the same place (TEMA or TEMS) and follow the same summarization and pruning schedule.

As stated earlier, it is important for you to develop a way to organize how managed systems are assigned to historical distributions. When you ⬀ **Create another collection**, the managed system distribution from the original collection definition is not copied from the original collection. This is to help ensure that you do not assign the same managed systems to other collection definitions for the attribute group. If you do, historical data samples will be taken for each collection definition, a symptom of which is the occurrence of duplicate rows in a table view. Group functions, such as averages, will have skewed results also. If this happens, use the CLI `tacmd`

histListCollections (see *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/ic/landing_cmdref.htm)) to list the historical collection distributions, then remove the managed system reference from all but one of the collections for the attribute group.

Historical attribute groups

Some attributes groups, such as Situation Status and Windows Event Log, are historical in nature and show all their entries without your having to specify a time span. You do not need to configure history collection for these attribute groups unless you want to roll off to a data warehouse or limit the data reported.

Configurable elements of collections for an attribute group and their dependencies

Summarization and pruning is configured by attribute group; not by the collection definitions, and can be configured independently of collection definitions.

Collection interval, the frequency that data samples are taken at the managed system and saved to the short-term history file, can be different for each collection definition.

Collection location for the short-term history file can be at the managed system or the monitoring server it connects to. There is one short-term history file for an attribute group per managed system; whatever was set for the first collection is used by all collection definitions for that attribute group.

Warehouse interval, for controlling how often historical data is copied from the short-term history file to the data warehouse, can be set differently for each collection definition.

Distribute to ⦿ Managed System (Agent) or ◯ **Managing System (TEMS)** can be set differently for each collection definition.

When this is set to *Managing System*, distribution is to all managed systems connected to the monitoring server; the ability to selectively distribute the collection on managed systems is disabled.

When this is set to *Managed System*, you have the option of assigning managed systems to the individual collection definition in the Distribution tab or to a historical configuration group that the collection is a member of or both. And managed systems can be assigned individually or as a managed system group.

Filter criteria for preventing the collection of unwanted data is configurable for each collection definition. Applying filters to historical data collection can help reduce network traffic and wasted disk space and can improve summarization and pruning performance.

Be aware that filtered data collection can affect the results of trending calculations that are performed with chart baselining and situation modeling and of query-based views that include historical data. For example, a filter to collect only rows where the process uses 80% or more CPU means that a calculation of average values will be only of values 80% and higher, rather than of all values.

Creating a historical collection

A historical data collection specifies the attribute group to collect data from, where to store the historical data, and other information such as the collection frequency and distribution. Create a historical collection definition for every attribute group that you want to collect historical data for. You can then retrieve the historical data into query-based views.

Before you begin

Your user ID must have Configure History permission to open the History Collection Configuration window. If you do not have this permission, you will not see the menu item or tool for historical configuration.

The CCC Logs apply to all managed systems: Agent Operations Log, ITM Audit, and Universal Messages. The Policy_Status and System_Status attributes do not record historical data. Typical environments do not use the EIB Change Log and Situation Status Log attribute groups.

About this task

Complete these steps for each attribute group that you want to collect historical data from on specified managed systems or managed system groups or on all the managed systems that connect to a Tivoli Enterprise Monitoring Server:

Procedure

1. Click **History Configuration** to open the History Collection Configuration window.
2. Click **Create a new collection**. If you first click a Monitored Application (or right-click and click), **Monitored Application** will be selected for you.
3. Enter a **Name** of up to 256 bytes. A short name is also given to the collection and is shown in the middle section of the status bar.
4. Optional: Enter a **Description** for the collection, up to 64 bytes.
5. Select a **Monitored Application** from the list.
6. Select an **Attribute Group** from the list. Not all the product attribute groups necessarily display: only those that are appropriate for historical collection and reporting.
7. Click **OK** to open the configuration tabs for the collection. The branch of the monitored applications expands to show the new collection. At this point, it has not been distributed to the managed systems.
8. Complete the fields in the **Basic** tab:
   a. **Collection Interval** is the frequency of data transmission to the short-term history file on the computer where the data is saved (Tivoli Enterprise Monitoring Agent or Tivoli Enterprise Monitoring Server). The options are every 1, 5, 15, or 30 minutes, every hour, or once per day. The default interval is 15 minutes.
      The shorter the interval, the faster and larger the history file grows. This can overload the database, warehouse proxy, and summarization and pruning agent. For example, if you set a 1-minute collection interval for Process data, expect the summarization and pruning for that attribute group to take a long time. Such a short interval should be enabled for an attribute group only if it is critical in your work.
   b. **Collection Location** is where the short-term historical data file resides: at the TEMA (Tivoli Enterprise Monitoring Agent) or the TEMS (Tivoli

Enterprise Monitoring Server). The default location is TEMA, which minimizes the performance impact on the monitoring server from historical data management.

If you are filtering the historical collection (step 9), collection at the TEMA also gives you more space for composing the filter. The formula example at the end of this topic takes up 63% of the formula storage when collection is set to TEMS; and 31% when collection is set to TEMA..
Note, however, that TEMS might be a better choice for certain environments. Also, the OMEGAMON® XE on z/OS product requires that the data be stored at the monitoring server.

c. **Warehouse Interval** determines whether the collected data is warehoused and how often. The options are 1 minute, 5 minutes, 15 minutes, 30 minutes, 1 hour, 12 hours, 1 day, or Off.

A more frequent warehousing interval enables quicker availability of warehoused data for retrieval. Shorter intervals cause some additional processing to check for and transmit newly collected data; and there are more frequent elevated levels of transmission activity, but for shorter durations. Regardless of the warehouse interval, the most recent 24 hours are always available in the short-term history files. If you use Tivoli Common Reporting, note that queries are to the data warehouse only.

9. If you want to restrict the collection of samples to only those that meet certain conditions, click the **Filter** tab and compose the formula:

a. Click **Add Attributes**.

b. Select the attribute to filter on and click **OK**. You can use Ctrl + click to select multiple attributes, Shift + click to select all entries between the attribute previously selected and this one, or **Select All** to select every attribute.

c. For each attribute selected, click inside the cell under the column heading and compose an expression $\boxed{v}\boxed{==}\boxed{\quad}$ consisting of a function, a relational operator and a test value:

To change the function, click $\boxed{v}$ Value (or Compare date and time for a time attribute) and select one from the list. See Appendix A, "Formula functions," on page 415 for the syntax, description, and an example of each function.

To change the relational operator, click $\boxed{==}$ equal and select another operator from the list.

Click inside the $\boxed{\quad}$ text box and enter or select the value.

d. Keep multiple expressions in the same row if they must all be met (Boolean AND logic) and on separate rows if any of them might be met (Boolean OR logic) for the sample to be saved in the short-term history file.

e. Select the same attribute twice if you want to collect a range, such as > 1 AND < 90. Use **Add Attributes** to select the attribute again. (See the second example at the end of this topic.)

f. Confirm that the **Formula capacity** counter is below 100%. Setting the **Collection Location** (step 8 on page 225) to **TEMA** gives the most space.

10. Click **Apply** to save the collection.

Results

The collection name appears in the monitored application branch of the tree with an icon to indicate that it is not distributed to managed systems: ▉ .

Example

This is a historical collection called "ProcessInformation". The settings in the **Basic** tab use the defaults: **Collection Interval** of 5 minutes; **Collection Location** at the TEMA; and **Warehouse Interval** of 15 minutes. The **Filter** tab has a formula written to collect the sample if one of these two conditions is met: the process name is *VirusScan*; OR the process name is *java* and is taking more than *50%* of the processor time on the *North012* Windows system, the *South007* Windows system, or both.

	Process Name	% Processor Time	Server Name
1	==VirusScan		
2	==java	>50	('Primary:North012:NT', 'Primary:South007:NT')

The IN operator was selected for entering the server list, as you can see when viewing the formula in :

```
(Process Name==VirusScan) OR (Process Name==java AND %Processor Time >50
AND Server Name *IN ('Primary:North012:NT','Primary:South007:NT'))
```

The following example shows a filter formula with a repeating attribute:

	% Processor Time	% Processor Time
1	>40	<60

What to do next

After creating a new collection definition, you must distribute it to the managed systems where you want to take data samples. You can assign the managed systems to distribute the collection to in the Distribution tab or assign the collection to a historical configuration group that has a distribution. If you are not ready to distribute the collection, you can do so at a later time.

If you added a filter to restrict the collection data samples to samples that meet the criteria, test the formula to make sure that the intended data is being collected. It is possible to enter an incorrect value, which might result in no data being collected. To test the filter formula, wait until enough historical data has been collected to show in a historical table view of the attributes. Then set a time span for the table view. If no rows are retrieved to the table view, review the filter formula for typographical or value errors.

Related concepts:

Managed system groups
A managed system group is a named collection of managed systems. When distributing a situation, policy, or historical collection configuration to managed systems, or when assigning managed systems to a custom Navigator view, you can specify a managed system group (or groups).

Distributing a historical collection

Distribute a historical collection definition to the managed systems that you want to save data samples from. Data collection begins as soon as you save the distribution. Edit the distribution list whenever you want to start or stop historical sampling on certain managed systems or managed system groups.

Before you begin

After you create a new historical collection definition or create another, the next step is to select the distribution method (Managed System or Managing System) and assign managed systems to the distribution list.

About this task

Take these steps to create or edit the managed system distribution for a historical collection definition:

Procedure

1. If the History Collection Configuration window is not open, click ⬛ **History Configuration**, click the **Monitored Application**, then click the ⬛ historical collection definition that you want to distribute.

2. Click the **Distribution** tab and select the distribution method:

   - ◉ **Managed System (Agent)** to select managed systems and managed system groups individually. You can also add the collection definition to a historical configuration group to adopt the group's distribution.

   - ○ **Managing System (TEMS)** to select all the agents that connect to this monitoring server. This method makes the collection ineligible for inclusion in a historical configuration group.

   In the following example, if you chose distribution by *managed system (agent)*, you could select any of the agents to distribute the collection to; if you chose distribution by *managing system (TEMS)*, selecting the East Tivoli Enterprise Monitoring Server distributes the collection to the 1, 2, and 4 agents and you cannot remove any of them or add agents from the West monitoring server.

   Available Systems shows the monitoring agents or monitoring servers for assignment and **Available Managed System Groups** shows the available managed system groups. You can click **Edit Managed System Groups** to open the Object Group editor for creating and editing managed system groups.

3. If you selected ◉ **Managing System (TEMS)**, select the monitoring server whose managed systems you want to start collecting data for, and ◀ move to the **Start collection on** list.

4. If you selected ◉ **Managed System (Agent)**, distribute to the managed systems where you want to start data collection:

   - Assign managed systems to the collection: Select from the **Available Systems** or **Available Managed System Groups** list, then ◀ move to the **Start**

collection on list. Distribution to an individual historical collection is useful when the managed system assignments vary for each collection definition.

- Assign managed systems to a historical configuration group for the collection definition. Distribution to a historical group is useful when the distribution is the same for multiple collection definitions.

  a. Click **Group**.

  b. Click **Create a new object group**; or click **Update an existing object group** if a historical group has already been created that this collection can become a member of.

  c. In the **Members** tab, select a historical collection from the list of **Available Members** and ◀ move it to the **Assigned Members** list.

  d. In the **Distributions** tab, select the managed systems and managed system groups to assign to the historical configuration group. Select a managed system group to see its constituent objects in the **Selected Assigned Object Members** or **Selected Available Object Members** list.

  e. Click **OK** to save the historical group and return to the Historical Collection Configuration window.

  The distribution from the historical configuration group (or groups) that the collection belongs to is shown in the **Start collection on** list; the entries are dimmed to indicate that they were assigned through a historical configuration group and cannot be changed from here. You can right-click a dimmed entry and click **Show Groups** to see the historical groups it belongs to.

5. Click **Apply** to save the collection definition and start data collection on the managed systems.

Results

The collection shows in the tree as 🖼 distributed to indicate that data collection is taking place on the assigned managed systems.

What to do next

Click the 🔘 monitored application name in the tree and, in the right frame, find the attribute group for the collection that you just configured. The attribute group is prefixed with a 🏃 **Data collection started** icon.

If the collection definition has a warehouse interval, you need to configure summarization and pruning for the attribute group if you have not already done so.

After historical data samples have been saved to the short-term history files, you can create query-based views and select a time span of historical data to display.

After data collection has begun, the **Distribution** method radio buttons are disabled. If you want to enable them to change the distribution method, you must first remove the current distributions as described in the topic, "Changing the Distribution Method".

Related tasks:

Creating a managed system group
Organize managed systems into named groups that can be assigned to the distribution of situations, policies, historical collections, queries, and dynamic Navigator items in custom Navigator views.

Configuring summarization and pruning for attribute groups

Configure summarization and pruning for the Tivoli Data Warehouse to aggregate data and keep the database size at a manageable level.

Before you begin

The summarization and pruning agent must be installed, configured, and started as described in the *IBM Tivoli Monitoring Installation and Setup Guide*.

Your user ID must have Configure History permission to open the History Collection Configuration window. If you do not have this permission, you will not see the menu item or tool for historical configuration.

About this task

Although summarization and pruning is not mandatory for warehoused data, it keeps the database from growing to an unwieldy size and minimizes the amount of data that gets retrieved to the Tivoli Enterprise Portal. Even if data collection for an attribute group has not been configured, you can set up summarization and pruning. If no collections have been created and distributed for an attribute group, no data goes to the warehouse, and summarization and pruning does not take place.

Procedure

1. If the History Collection Configuration window is not open, click ☑ **History Configuration**.
2. Select a **Monitored Application** from the tree.
3. Review the attribute groups in the table. If summarization and pruning has already been configured for an attribute group, the values will be shown in the summarization and pruning cells. Collapse the tree, drag the borders, or scroll the table right to see all the cells.
4. Select one or more attribute groups to configure. You can select multiple groups with Ctrl+click, or Shift+click to select all groups from the first one selected to this point. The settings of the first group selected continue to display, regardless of the settings in any of the other selected groups. This enables you to adjust the configuration controls once and apply the same settings to all selected attribute groups. Use the **Clear all** button if you want to clear all the fields and start over.
5. In the **Summarization** area, select the check box for every time period to be aggregated: Yearly, Quarterly, Monthly, Weekly, Daily, and Hourly.
6. In the **Pruning** area, select the check box for every time period to be pruned: Yearly, Quarterly, Monthly, Weekly, Daily, and Hourly. If you also want to keep the original data samples, select the **Detailed data** check box. In the corresponding fields, specify the number of days, months, or years to keep the data.

7. Click **Apply** to save the configuration for the attribute groups that were selected. The summarization and pruning cells for the attribute group are updated to reflect the new settings.

What to do next

The next time summarization and pruning takes place, the summarization and pruning agent applies the configuration to the long-term data stored in the data warehouse. Wait for the next scheduled time period to elapse before expecting to see any summarized data.

To configure granular summarization and pruning for managed system groups, you must use the tacmd Command Line Interface (CLI). For detailed instructions, see the *IBM Tivoli Monitoring Administrator's Guide*.

Creating another historical collection

After historical collection definitions have been created, use the History Collection Configuration window to create duplicates that you can customize for different collection intervals, warehousing intervals, and distribution.

Before you begin

Your user ID must have Configure History permission to open the History Collection Configuration window. If you do not have this permission, you will not see the menu item or tool for historical configuration.

About this task

Take these steps to copy a historical data collection:

Procedure

1. Click ⬚ **History Configuration** to open the History Collection Configuration window.
2. Click the ⬚ managed application, then the collection that you want to copy.
3. Click ⬚ **Create another collection**.
4. Enter a new name for the collection and, optionally, a description. The **Monitored Application** and **Attribute Group** fields are disabled.
5. Click **OK** to open the configuration tabs for the collection, make any changes to the intervals, and assign the managed systems to distribute the collection to.

What to do next

Other than the **Collection Location**, which uses what was specified for the original collection definition, configure the data collection just as you would a new collection. You can change the collection and warehouse intervals, and you must specify the distribution.

Creating a historical configuration group

Organize the historical collections that you have defined into named groups that can be managed collectively.

Before you begin

Historical configuration groups are made up of historical data collections that were defined in the History Collection Configuration window. Create the historical collections that you intend to include in the historical group before opening the Object Group editor.

This feature is available when the Tivoli Enterprise Portal Server and Tivoli Enterprise Monitoring Server are at Version 6.2.2 or later and the historical data collection that you want to add to the group is configured for **Managed System (Agent)** as the distribution method.

About this task

Take these steps to create a historical configuration group:

Procedure

1. Open the Object group editor from either of these locations:
   - Click ⊞ **Object group editor**, then click ⬛ **Historical configurations**.
   - If the History Collection Configuration window, is not open, click ⬛ **History Configuration**. Click a ⬤ monitored application, click an ▬ undistributed or ⬛ distributed historical collection definition, then click **Group**.
2. Click ⬚ **Create a new object group**.
3. Type a name for the historical group and, optionally, a description, then click **OK**. The new historical configuration group is displayed in the ⬛ **Historical configurations** branch.
4. In the **Members** tab, select a historical collection from the list of **Available Members** and ◂ move it to the **Assigned Members** list. You can select multiple collection definitions for an attribute group by holding down Ctrl while clicking; or use Shift+click to select all entries between this and the first selection. Historical collections that are configured for "Managing System (TEMS)" distribution are not displayed in the list.
5. If you have historical configuration groups available and want to assign a group to this one, click the **Member of** tab, then select a group and ◂ move it to the **Assigned Members Of** list.
6. In the **Distributions** tab, select the managed systems and managed system groups to assign to the historical configuration group. Select a managed system group to see its constituent objects in the **Selected Assigned Object Members** or **Selected Available Object Members** list.
7. Save your changes and either keep the editor open with **Apply** or exit with **OK**.

Results

The historical configuration group is now available to all users who log onto the same portal server. Historical data samples are collected to the short-term history files for the managed systems in the distribution.

When you next view the historical data collection in the History Collection Configuration window, the Distribution tab shows the historical configuration group in the **Start collection on** list. The group name is disabled (dimmed); you cannot remove it from the distribution. To do that, you would return to the Object group editor and remove the historical collection from the historical group.

Related tasks:

Creating a managed system group
Organize managed systems into named groups that can be assigned to the distribution of situations, policies, historical collections, queries, and dynamic Navigator items in custom Navigator views.

Stopping data collection

Stop data collection on one or more managed systems for a historical collection definition when you are no longer interested in gathering data on the managed system or would like to temporarily disable data collection.

Before you begin

Your user ID must have Configure History permission to open the History Collection Configuration window. If you do not have this permission, you will not see the tool for historical configuration.

About this task

The historical collections for an attribute group can be edited to stop data collection on some or all of the managed systems in the distribution list. The managed systems that remain in the distribution continue to collect historical data. Take these steps to stop data collection on a managed system or managed system group:

Procedure

1. Click ☑ **History Configuration** to open the History Collection Configuration window.
2. Click the ⊚ monitored application (agent type).
3. Click the ☑ collection that the managed system is assigned to.
4. For managed systems that are assigned to the historical collection:
   a. Click the **Distribution** tab to see the managed systems, managed system groups, and historical object groups that collection is occurring on.
   b. In the **Start collection on** list, click the managed systems or managed systems groups (or both) that you want to stop collection on and ▶ move them back to the **Available** list. You can also use Ctrl+click to select multiple items and Shift+click to select all items between the item last selected and this one.
   c. If this is a historical object group, the name is disabled as are any managed systems and managed system groups that are assigned to the distribution through the historical group. Right-click a name and click ⊞ **Show Groups** to see the historical configuration group it was distributed to, then complete the next step.
5. For managed systems that are assigned to a historical configuration group:
   a. Click **Group**.
   b. Click the ⊞ historical group in the ☑ **Historical configurations** branch.
   c. Click the **Distributions** tab.
   d. In the **Assigned managed systems** list, click the systems or groups (or both) that you want to stop collection on and ▶ move them back to the **Available** list.
6. To save your changes, click **Apply** to keep the window open or **OK** to close it.

Results

Historical data collection no longer occurs on the managed systems that you moved out of the **Start collection on** list for the historical collection definition or that you moved out of the **Assigned managed systems** list for the historical configuration group.

If data collection was stopped on all managed systems that were assigned to the historical collection, its icon changes to indicate that it is ▬ not distributed. Likewise, if data collection was stopped on all managed systems for an attribute group, the 🏃 icon that indicates data collection is started in the summarization and pruning table no longer displays.

What to do next

If you want to verify that you removed a managed system or managed system group from all historical collection definitions, you can use the `tacmd histlistcollections` command to list all the collections that are started for a managed system or managed system group.

For example, `histListCollections -m Primary:BRAZIL:NT` shows all historical collections that are started on the managed system called Primary:BRAZIL:NT. See the *IBM Tivoli Monitoring Command Reference* for more information.

If your intention is to stop data collection on all managed systems for this collection permanently, delete the historical collection definition; if your intention is to stop data collection for the entire attribute group, delete all the collection definitions that were created for it.

Starting data collection

You can start collection of historical data for attribute groups on selected managed systems and whenever you have managed systems that you want to add.

Before you begin

Your user ID must have Configure History permission to open the History Collection Configuration window. If you do not have this permission, you will not see the menu item or tool for historical configuration.

About this task

After you distribute the collection definitions to managed systems, data collection begins automatically. If data collection has been stopped for any of the managed systems in the distribution, complete these steps to resume collection. Complete these steps also if you want add new managed systems to the distribution.

Procedure

1. Click 🗒 **History Configuration** to open the History Collection Configuration window.
2. Click the ◉ monitored application (agent type).
3. In the expanded branch, click the ▬ undistributed or 🗒 distributed historical collection that you want to start on managed systems.

4. Assign managed systems to the collection or add the collection to a historical configuration group to use its distribution. See "Distributing a historical collection" on page 227 for details.

5. To save the revised distribution and begin data collection on the managed systems that you added, click **Apply** to keep the window open or **OK** to close it.

Changing the distribution method

Change the distribution method from **Managing System (TEMS)** to **Managed System (Agent)** to enable the historical collection definition to be added to a historical object group.

About this task

If you have upgraded to Tivoli Management Services Version 6.2.2 from a previous version, you will see a historical collection for each attribute group that was configured, and the distribution method will be as it was before: **Managing System (TEMS)**. Complete these steps if you want to change the distribution method for each historical collection to **Managed System (Agent)**.

If, instead, you want to change the distribution method from Managed System to Managing System, the procedure is the same except that it can be more involved if an entry in the **Start collection on** list is disabled. A disabled entry means that it was added to the distribution as part of a historical configuration object group. If any of the collection belongs to one or more historical groups, you must remove it from any historical group that it belongs to that has a distribution before you can remove it from the **Start collection on** list. The topic on "Stopping data collection" has instructions for removing managed systems from the distribution of a historical object group.

Procedure

1. If the History Collection Configuration window is not open, click ⬛ **History Configuration**, click the **Monitored Application**, then click the historical collection name to open the definition.

2. Click the **Distribution** tab.

3. Select all entries in the **Start collection on** list, and ▶ move them to the **Available** list. After clicking the first entry, you can Shift+click the last entry to select all of them.

4. Click **Apply**.

What to do next

You can now assign managed systems to the distribution here or add the historical definition to a historical object group to use the distribution of the group.

Removing a historical data collection

Use the History Collection Configuration window to remove any unnecessary historical data collection definitions.

Before you begin

Your user ID must have Configure History permission to open the History Collection Configuration window. If you do not have this permission, you will not see the menu item or tool for historical configuration.

About this task

Take these steps to remove a historical collection definition and thereby no longer collect data for the attribute group on the assigned managed systems. Be aware that if there are multiple collection definitions for the same attribute group, you must delete all of them before data collection will be discontinued for the attribute group.

Procedure

1. Click ▱ **History Configuration** to open the History Collection Configuration window.
2. Click the ⊙ monitored application (agent type).
3. In the expanded branch, select the historical collection that you want to unconfigure. You can use Ctrl + click to select multiple collections or Shift + click to select all between the one previously selected and this one.
4. Click ✖ **Delete collection**.
5. Confirm the deletion of the historical collection definition to permanently remove it.

Results

Data collection stops for the historical collection definition on all managed systems it was distributed to. The historical data collected up to this point continues to remain in the short-term history files and data warehouse.

Chapter 9. Navigator views for logical organization

Create custom Navigator views to organize your monitored environment into logical hierarchical groupings. You can also create Navigator views and build workspaces that are targeted for specific points of view, such as executive, operator, or subject matter expert for an application domain.

The Navigator editor enables you to create new Navigator views and build a hierarchical structure by creating new *child* items as branches in the tree and sharing items from a source Navigator view with the new (target) view. You can assign managed system groups to dynamically update Navigator items whenever the group is edited.

Custom Navigator view characteristics

Custom Navigator views and the Navigator Physical view have characteristics that distinguish them from one another. Understanding these distinctions and the customization options can help you design the most effective Navigator views.

Navigator Physical view

The Navigator Physical view is the default Navigator view. It shows your monitored enterprise as a physical mapping of managed systems by their node names and operating systems. The Navigator Physical view is a discovered view: as managed systems are added to the monitored environment, the Tivoli Enterprise Portal Server discovers them and displays them in the Navigator view.

As well as showing the physical hierarchy of your environment, the Navigator gives you visual indication of situation events that occur on managed systems associated with a Navigator item and organizes your workspaces.

You can create situations and associate them with the Navigator items where you want to see alert indicators and you can create workspaces and report data from the managed systems assigned to that branch of the tree. But you cannot edit the Navigator Physical view itself.

Custom Navigator views

You can create Navigator views for logical hierarchies, such as by location or business application. With custom Navigator views, you can create Navigator items and determine their managed system assignments. You can add Navigator items either by dragging and dropping them from a source Navigator view to the target view or by creating them manually.

The drag and drop method *shares* the item in the source and target Navigator views. When you share an item, all managed system assignments, workspaces, and situations associated with the source item are applied. Future changes to one item are applied to the other. A ⊕ **shared** icon is displayed over the source and target Navigator item icons. It provides a visual reminder of the relationship so that you do not unintentionally change or delete an item from one Navigator view that might affect another view. You cannot change the managed system assignments of Navigator items that were shared using the drag and drop method.

For Navigator items added with 🔧 **Create child item**, you control the management system assignments:

- Navigator items with no managed systems assigned cannot have situations associated with them, thus no alert indicators are displayed on the item except as part of the roll-up display of events. This might be what you want for a container (parent) Navigator item.

- Navigator items can have managed system assigned *statically*. When creating the item, you assign managed systems or managed system groups (or both) through the **Monitored Resources** tab of the Navigator item properties. Initially, static items are not shared with other items nor are situations associated with them.

- Navigator items can have managed systems assigned *dynamically*. When creating the item, you assign a managed system group (or groups) through the **Dynamic Items** tab. The Navigator item is then populated by all the managed systems that belong to that group. For example, you might have a Navigator item named "Far West" with "Arizona" and "California" managed system groups assigned as dynamic items. Any managed systems added or removed from those groups are reflected in the "Far West" branch. The items at the 🔘 agent and 🖥 attribute level are shared from the Navigator Physical view.

  📁 Support
     ⊟ 📁 Far West
        ⊞ ⊞ Arizona
        ⊟ ⊞ California
           ⊞ 📟 JOSE95101
           ⊞ 📟 LEANDRO94579
           ⊟ 📟 MARINO91108
              ⊟ 🔘 UNIX OS
                 ⊟ 🖥 Disk Usage
                 ⊟ 🖥 File Information
                 ⊟ 🖥 Network
                 ⊟ 🖥 ...

  Options for dynamic members enable you to hide any or all of these items: ⊞ group level, 📟 node level, or 🖥 attribute level. For example, the Support branch looks like this when all options are disabled:

  📁 Support
     ⊟ 📁 Far West
        ⊟ 🔘 UNIX OS

- You can assign managed systems and managed system groups to a new Navigator item, including the assignment of managed system groups as dynamic members.

Creating a new Navigator view

The Navigator Physical view is organized by managed systems. Create custom Navigator views to organize your enterprise differently, such as by geography or user role.

Before you begin

Your user ID must have Modify permission enabled for the *Custom Navigator Views* feature.

Procedure

1. In the Navigator toolbar, click ☑ **Edit Navigator View**. The Edit Navigator View window opens with the Navigator Physical view as the Source View on the right and the Logical view or most recently opened Navigator view as the Target View for editing.

2. Click ⬚ **Create New Navigator View**. A window opens for you to enter the Navigator view identity.

3. Type a name for the Navigator and, optionally, a description, then click **OK**. After you click **OK**, the Navigator view name is displayed as the ⬚ enterprise item. If, instead, you get a message that the name already exists, enter a different name. Because Navigator views are assigned to user IDs, you might not be able to see Navigator views that other users have created and it is therefore possible to enter duplicate names.

4. Take any of the following steps to populate the target Navigator view:

   • To create a new Navigator item, select or right-click the parent item and click ⬚ **Create Child Item**. See "Adding a child item" on page 240.

   • To share a Navigator item, drag it from the source view on the right to the target on the left and drop it on the item it should follow. (After selecting the first item, use Ctrl+click to select other items or Shift+click to select all items between the first selection and this one.) If you want to share Navigator items from a different source view than the one displaying, select the view to use from the **Source View** ▾ list.

   • To assign managed systems to a Navigator item, right-click the item and click ⬚ **Properties**. The managed system assignments make the item eligible for situation association and are referenced in query-based views.

     – Click the **Monitored resources** tab for assignments to this item.

     – Click the **Dynamic items** tab for dynamic child assignments to this item. See "Navigator item properties" on page 246.

     – Click the **Symbols** tab to define symbol name/value pairs associated with a physical or logical navigator node. See "Adding or editing a symbol" on page 249.

     ⚑ A best practice is to assign managed systems to a Navigator item from the **Monitored resources** tab or **Dynamic items**, but not both.

   • To find out if the Navigator item is assigned to other Navigator views or assigned multiple times to this or other Navigator views, right-click the item and click ⬚ **Show Navigator List**. A number in parentheses after the view name indicates that there are multiple instances of this Navigator item in the Navigator view. You can also switch to another Navigator view in the list.

   • To rename an item, right-click it and click ⬚ **Properties**; to rename the view, click ⬚ **Properties**. Edit the name and description as needed.

   • To delete a Navigator item, select it and click ⬚ **Delete Item**.

5. Click **Apply** to save your changes to the Navigator view and keep the editor open; or click **OK** to save your changes and close the window.

Results

The new Navigator view is displayed in the Navigator view list. Shared Navigator items and dynamic Navigator items show alert indicators for any open events for associated situations. Shared Navigator items also share the workspaces from the source item; all other Navigator items have undefined workspaces.

Note: In the Navigator Physical view, Tivoli Enterprise Portal indicates an agent has gone offline by dimming the agent branch for that system. This indication occurs in custom Navigator views as well. However, if you share an attribute-level item from the Navigator Physical view with a custom view without sharing its parent agent-level item, you receive no indication when the agent is offline: The attribute item appears normal but when you open its workspace, any chart or table views show no information.

What to do next

- Open the new Navigator view and review the items to confirm that they reflect the managed system and dynamic member assignments that you created.
- Create workspaces for Navigator items that have none defined.
- If you want the new Navigator view to be available to other users, assign it in the Navigator views tab of the Administer Users window. You can also assign a user to a specific branch in the Navigator, called the *Assigned Root*, or grant the user access to a specific set of agents, called *Allowed Applications*.
- For items that are not shared, alert indicators for situation events do not automatically show on the Navigator item for every assigned managed system that also has a situation distributed to it. If you want alert indicators, you must associate the situation with a Navigator item.
- For items with static managed system assignments, any new managed systems that come online that you want in the view must be added through the Navigator item properties or by dragging and dropping from the Navigator Physical view to the target view in the Navigator editor.

Related tasks:

"Adding a child item"
Use ⁂ **Create Child Item** in the Navigator editor to add a Navigator item within the selected branch of the custom Navigator view.

"Editing Navigator item properties" on page 243
Open the Navigator item properties to see and edit the name, description, and managed system assignments.

"Deleting a Navigator view or item" on page 244
Use the Navigator editor to delete a Navigator view or any of its Navigator items.

"Associating a situation" on page 287
If you access the Situation editor through the toolbar to create situations, no alert will be visible in the Navigator when the situation becomes true until you associate the situations with a Navigator item.

Adding a child item

Use ⁂ **Create Child Item** in the Navigator editor to add a Navigator item within the selected branch of the custom Navigator view.

About this task

Use the following steps to add a child item to a custom Navigator view:

Procedure

1. If the Navigator editor is not open, click 📝 **Edit Navigator View** in the Navigator toolbar.
2. If the Navigator view that you want to work with is not displayed, select it from the **Target View** ⏷ list.

3. Select the parent item in the target view. The item has a border around it when it is selected.

4. Click ⃰ **Create Child Item**.

5. In the Create New Navigator Item window, type a name and description for the item. The name is displayed in the Navigator view; the description is displayed only in this window.

6. Make static or dynamic managed system assignments: Select the managed systems or managed system groups whose situations you want to associate with the Navigator item, and click ◄ to move them to the **Assigned** list. Queries for workspace views are applied to the assigned managed systems by default. Alert indicators are displayed for events that are opened for situations associated with the Navigator item.

   - For assignments to this Navigator item, click the ▣ **Monitored resources** tab, select the managed systems or managed system groups whose situations you want to associate with the Navigator item, and ◄ move them to the **Assigned** list.

   - For dynamic child assignments to this Navigator item, click the ⊞ **Dynamic items** tab, select the managed system groups to assign to the Navigator item, and ◄ move them to the **Assigned Members** list. Click ▦ **Options** to change how branch dynamic items are displayed (described in "Assigned member options" on page 249).

   - If you are assigning managed system groups to the Navigator item, you can click **Edit Managed System Groups** to open the Object Group editor: select a managed system group to see and edit the membership; select a managed system type (select ▣ **All Managed Systems** if you want to assign managed systems from different managed system types) and click ▯ **Create New Group** or ▯ **Create Another Group** to add a new managed system group (see "Creating a managed system group" on page 387).

   - Click the ⓧ **Symbols** tab to define symbol name/value pairs associated with a physical or logical navigator node. See "Adding or editing a symbol" on page 249.

   ⚐ A best practice is to assign managed systems to a Navigator item from the **Monitored resources** tab or **Dynamic items**, but not both.

7. Click **OK**.

Results

- The child item is displayed within the parent item hierarchy, below any other child items, with this icon: ▢ .

- A ⊕ shared icon is displayed over the Navigator items that have been shared. This indicator serves as a reminder so that you do not unintentionally change or delete an item from one Navigator view that might affect another view.

- A dimmed **<Dynamic members>** child item indicates that managed systems are assigned dynamically.

What to do next

If the child item is in the wrong branch, you can drag it to the desired location to create a duplicate item there, then delete the original item.

Editing a Navigator view

Use the Navigator editor to edit the Navigator view and the managed systems or dynamic members assigned to the items.

Procedure

1. Select the Navigator view you want to edit from the ⬇ list in the Navigator toolbar. You cannot edit the Navigator Physical view.
2. Click ✎ **Edit Navigator View**.
3. If the Navigator view you want to edit is not displayed as the target, select it from the **Target View** ⬇ list. Likewise, you can select a different source from the **Source View** list.
4. Take any of the following steps to populate the target Navigator view:
   - To create a new Navigator item, select or right-click the parent item and click ⛁ **Create Child Item**. See "Adding a child item" on page 240.
   - To share a Navigator item, drag it from the source view on the right to the target on the left and drop it on the item it should follow. (After selecting the first item, use Ctrl+click to select other items or Shift+click to select all items between the first selection and this one.) If you want to share Navigator items from a different source view than the one displaying, select the view to use from the **Source View** ⬇ list.
   - To assign managed systems to a Navigator item, right-click the item and click ⬚ **Properties**. The managed system assignments make the item eligible for situation association and are referenced in query-based views.
     – Click the **Monitored resources** tab for assignments to this item.
     – Click the **Dynamic items** tab for dynamic child assignments to this item. See "Navigator item properties" on page 246.
     – Click the **Symbols** tab to define symbol name/value pairs associated with a physical or logical navigator node. See "Adding or editing a symbol" on page 249.

     🎗 A best practice is to assign managed systems to a Navigator item from the **Monitored resources** tab or **Dynamic items**, but not both.
   - To find out if the Navigator item is assigned to other Navigator views or assigned multiple times to this or other Navigator views, right-click the item and click ⛁ **Show Navigator List**. A number in parentheses after the view name indicates that there are multiple instances of this Navigator item in the Navigator view. You can also switch to another Navigator view in the list.
   - To rename an item, right-click it and click ⬚ **Properties**; to rename the view, click ⬚ **Properties**. Edit the name and description as needed.
   - To delete a Navigator item, select it and click ⛌ **Delete Item**.
5. Click **Apply** to save your changes to the Navigator view and keep the editor open; or click **OK** to save your changes and close the window.

What to do next

Click the activated ⚙ **Navigator updates pending** tool in the Navigator toolbar to refresh the Navigator view.

"Editing Navigator item properties"
Open the Navigator item properties to see and edit the name, description, and
managed system assignments.

"Renaming a Navigator view or item" on page 244
Use the Navigator editor to rename a Navigator view or any of its Navigator items
or to edit the description.

"Deleting a Navigator view or item" on page 244
Use the Navigator editor to delete a Navigator view or any of its Navigator items.

Editing Navigator item properties

Open the Navigator item properties to see and edit the name, description, and
managed system assignments.

Before you begin

You can open the Navigator item properties in the Navigator view or in the
Navigator editor. When the item properties are opened from the Navigator view,
managed system assignments cannot be edited unless the Navigator item is not
shared. (A shared item is indicated by a ✛ shared icon over the Navigator item.)

You cannot edit the Navigator Physical view managed system assignments.

Procedure

1. Open the Navigator item properties from one of the following places:
   - Right-click the Navigator item and click 🔠 **Properties**.
   - In the Navigator toolbar, click 📝 **Edit Navigator View**. Right-click the
     Navigator item and click 🔠 **Properties**.
2. Edit the name and description as needed. The name is shown in the Navigator
   view and the description displays when you move the mouse pointer over the
   name in the tabbed list at the bottom of the view.
3. Edit the managed system assignments as needed. The Assigned list has no
   effect on the availability of queries for the workspace, only the event indicators
   for situations. If you have no managed systems assigned to this Navigator item,
   no events are displayed for it unless they are part of a roll-up display of events,
   and the Situation editor will not be available from the pop-up menu.
   - 🔘 **Monitored resources** shows the managed systems that were assigned
     through this tab or by shared Navigator items. If the Assigned list is
     populated but unavailable (dimmed), the assignments were created by
     sharing items from the Source Navigator view with the Target view.
   - ⊞ **Dynamic items** shows the managed system groups that were assigned
     through this tab. Use the tab to assign managed system groups to the
     Navigator item as *dynamic members* (child items) so that the Navigator item
     and the items it contains are updated automatically when managed systems
     are added to or removed from the managed system groups. Select a member
     and click 🔠 **Options** to change the dynamic Navigator item levels that are
     displayed in the branch.
   - 📄 **Symbols** shows the symbol name/value pairs associated with a physical
     or logical navigator node. You can also define Take Action commands that
     reference these symbol names. See "Adding or editing a symbol" on page
     249.

What to do next

Click the activated ⚙ **Navigator updates pending** tool in the Navigator toolbar to refresh the Navigator view.

Related tasks:

"Editing a Navigator view" on page 241
Use the Navigator editor to edit the Navigator view and the managed systems or dynamic members assigned to the items.

Renaming a Navigator view or item

Use the Navigator editor to rename a Navigator view or any of its Navigator items or to edit the description.

Procedure

1. In the Navigator toolbar, click ✏ **Edit Navigator View**.
2. If the Navigator view whose item you want to rename is not displaying, select it from the **Target View** ▾ list.
3. To rename the Navigator view, click ▦ **Properties** in the **Target View** toolbar. Type the new name in the **Name** field, add or change the description in the **Description** field, and click **OK**.
4. To rename a Navigator item, right-click the item and click ▦ **Properties**. Type the new name in the **Name** field, add or change the description in the **Description** field, and click **OK**.

Results

The Navigator view description is displayed in a flyover message box when you move the mouse pointer over its ⚙ tab; the Navigator item description is displayed only in this Properties window.

Related tasks:

"Editing a Navigator view" on page 241
Use the Navigator editor to edit the Navigator view and the managed systems or dynamic members assigned to the items.

Deleting a Navigator view or item

Use the Navigator editor to delete a Navigator view or any of its Navigator items.

Before you begin

Be careful when deleting Navigator items. A ⊕ shared icon is displayed over the source and target Navigator item icons to indicate a relationship with another item. This indicator serves as a reminder of the dependency so that you do not unintentionally change or delete an item from one Navigator view that might affect another view. To see the Navigator views this item is assigned to, right-click the item and click ⚙ **Show Navigator List**.

About this task

Complete the following steps to delete a Navigator view or item:

Procedure

1. In the Navigator toolbar, click ✏ **Edit Navigator View**.

2. If the Navigator view you want to delete or whose item you want to delete is not displayed, select it from the **Target View** ▾ list.
3. Take one of the following actions:
   - To delete the Navigator view, click ✖ **Delete Navigator View**. When a message asks if you are sure, click **Yes** or click **No** to cancel.
   - To delete an item, select it and click ✖ **Delete Item**. When a message asks if you are sure, click **Yes** or click **No** to cancel.

What to do next

- If you delete the current Navigator view (the one showing when you opened the Navigator editor), ⟳ **Navigator updates pending** is displayed in the Navigator view toolbar. Click this tool to remove the deleted Navigator view and open the next Navigator view.
- Deleted Navigator items are not removed from any graphic views that you have created in workspaces for the Navigator view. You need to delete the items manually from the view itself.
- When an agent is removed from the managed network, ⟳ **Navigator updates pending** is displayed in the Navigator Physical view toolbar for you to refresh the view. This also happens for custom Navigator views with dynamic Navigator item members. For Navigator items with no dynamic members (child items), you must delete the agent's Navigator item to remove it from the display.

Related tasks:

"Editing a Navigator view" on page 241
Use the Navigator editor to edit the Navigator view and the managed systems or dynamic members assigned to the items.

Navigator editor and properties

The Navigator editor is used to create and edit custom Navigator views and has other dialogs to support customization. Two of these dialogs, **Navigator item properties** and **Navigator List** are available from two places: the current Navigator view and the Navigator editor.

Edit Navigator view

Click ✎ **Edit Navigator View** in the Navigator toolbar to open the Navigator editor for creating and maintaining custom Navigator views.

The window has two frames:

Target View
> The **Target View** shows the name of the Navigator view you are editing in the left frame. When multiple Navigator views are available, you can select a different target from the ▾ list.

> ⟳ **Refresh now**
>> This refreshes the display of the target Navigator levels.

> ⊞ **Properties**
>> This opens the Properties window of the target Navigator view for you to see and edit the Navigator name and description.

>> You can also right-click a Navigator item and click ⊞ **Properties** to see and edit the monitored resources or the dynamic items.

> ⬑ **Create New Navigator View**
>> This opens the Create New Navigator View window for adding a

new Navigator view. After you have created the view, you can create new child items and assign managed systems (statically or dynamically) or drag items that you want to include in the view from the source Navigator view and rename them as needed.

Delete Navigator View
This deletes the current target Navigator view.

Create Child Item
Creates a new Navigator item below the one selected in the target view.

Delete Item
Deletes the selected Navigator item. Before deleting an item, right-click and click **Show Navigator List** to see if the item is repeated in this Navigator or is in other Navigator views.

Expand
Expand a Navigator item to see the child items. A dimmed **<Dynamic members>** child item indicates that managed systems are assigned dynamically.

A shared icon is displayed over the Navigator items that have been shared. This indicator serves as a reminder so that you do not unintentionally change or delete an item from one Navigator view that might affect another view.

Right frame
Source View shows the name of the original Navigator view in the right frame. Use this view to copy Navigator items to the target view. If you are editing the target view and not copying any items from this view, you can ignore what is in the right frame. When multiple Navigator views are available, you can select a different source from the ▾ list.

Refresh now
Refreshes the display of the source Navigator levels.

Properties
Opens the Properties window of the originating Navigator view for you to see and edit the name and description. You cannot change the properties of the Navigator Physical view.

Navigator item properties

When you create a child item in the Navigator editor, this window opens for you to enter the identity and assign managed systems. You can also open the Navigator Item properties in the Navigator editor or Navigator view by right-clicking a Navigator item and clicking **Properties**.

The item **Name** is displayed in the Navigator view. Enter a **Description** if you want to provide more detailed information about the item and its usage. See and change the managed system assignments in the Monitored resources, Dynamic items, and Symbols tabs:

Monitored resources
The Monitored resources tab shows the managed systems and managed system groups that were assigned to the Navigator item and any that are available for assignment. When the items are not selectable (dimmed), it means that the managed system assignments were shared with the Navigator Physical view.

Assigned

Shows the managed systems and managed system groups that are assigned to this Navigator item. To unassign a managed system or managed system group from the Navigator item, select it and ▸ move it to the **Available Managed Systems** or **Available Managed System Groups** list.

Available Managed Systems

Shows the managed systems that can be assigned to the Navigator item. To add a managed system to the Navigator item, click the name and move it to the **Assigned** list. After selecting a managed system, you can select multiples by holding down the Ctrl key while selecting others individually or by holding down the Shift key to select everything between two items.

Available Managed System Groups

Named groups of managed systems. To assign a managed system group to the Navigator item, click the name and ◂ move it to the **Available Managed Systems** list.

Click **Edit Managed System Group** if you want to open the Object group editor for creating and modifying managed system groups.

Example

After assigning managed systems to the "Eastern Region" item in the ◉ **Monitored resources** tab, this is how the tree looks in the Navigator editor and in the Tivoli Enterprise Portal window after closing the editor:

📁 Shipping
   📁 Eastern Region

⊞ **Dynamic items**

The Dynamic items tab is available for custom Navigator view items; it does not show for Navigator Physical view items. Use this tab to assign managed system groups to the Navigator item as *Dynamic members* (child items) and to control how the branch is displayed.

Assigned Members

Shows the managed system groups that are assigned as dynamic members (child items) of this Navigator item. Select a member and click ▦ **Options** to change how branch dynamic items are displayed, described in "Assigned member options" on page 249.

Selected Assigned Objects Members

Shows the managed systems and managed system groups that comprise a managed system group. ⊞ Expand the list to see the component managed systems and managed system groups of the group selected in the **Assigned Members** list.

Available Members

Shows the managed system groups that can be assigned as dynamic members (child items) of the Navigator item. To add a managed system group, click the name and ◂ move it to the **Assigned Members** list. After selecting a group, you can select multiples by holding down the Ctrl key while selecting others individually, or using the Shift key to select everything between.

Selected Available Objects Members

Shows the managed systems and managed system groups that comprise a managed system group. ⊞ Expand the list to see the

component managed systems and managed system groups of the group selected in the **Available Members** list.

Examples

Navigator editor Target View after assigning "Vermont" as a dynamic member to the "Eastern Region" Navigator item:

📁 Shipping
    ⊟ 📁 Eastern Region
          📁 <Dynamic members>

After no further changes were made in the Navigator editor, the Navigator view looks like this:

📁 Shipping
    ⊟ 📁 Eastern Region
        ⊟ ▦ Vermont
            ⊞ 🖳 BOB123
            ⊞ 🖳 PAT123
            ⊟ 🖳 SUE123
                ⊟ ◉ Windows OS
                    ⊟ 🖥 Disk
                    ⊟ 🖥 Enterprise
                    ⊟ 🖥 Memory
                    ⊟ 🖥 ...

The view above reflects the default settings for Assigned member options. You can deselect any of the options to reduce the number of dynamic items in the Navigator view. For example, with only ☑ **Include host-level items associated with managed systems** selected, the view looks like this:

📁 Shipping
    ⊟ 📁 Eastern Region
        ⊟ 🖳 BOB123
            ◉ Windows OS
        ⊟ 🖳 PAT123
            ◉ Windows OS
        ⊟ 🖳 SUE123
            ◉ Windows OS

📑 Symbols

The Symbols tab can be used to define symbol name/value pairs that are associated with a physical or logical navigator node. Click **Add** to create a new symbol name/value pair. Select a row to **Edit** or **Remove** the pair.

Additionally, you can define Take Action commands that reference these symbol names. When the Take Action command is issued from the navigator node, the symbol references are replaced by the values that are associated with the node. Use the **Take Action** > ⚡ **Create or Edit** dialog to define the command template, prefixing symbol name references with &. For example, to reference an Item Symbol named aSymbol, type &aSymbol at the location in the command template where you want the substitution to occur.

Name Specify the symbol name. The name must not include any spaces or special characters other than a period (.) and an underscore (_).

Value Specify the symbol value. If you want to include spaces, use enclosing quotes.

Symbol values can be encoded as Link Wizard expressions that are evaluated before the value is substituted into the Take Action command. You can edit and save the value as necessary to achieve the desired result.

For more information, see "Adding or editing a symbol."

Adding or editing a symbol

Define symbol name/value pairs associated with a physical or logical navigator node.

Name Specify the symbol name. The name must not include any spaces or special characters other than a period (.) and an underscore (_).

Value Specify the symbol value. If you want to include spaces, use enclosing quotes.

Symbol values can be encoded as Link Wizard expressions that are evaluated before the value is substituted into the Take Action command. You can edit and save the value as necessary to achieve the desired result.

The easiest way to compose a value expression is to launch the Link Wizard from the node. To open the Expression Editor dialog, create a dummy dynamic link, then navigate to the Target Filters page, and then select any of the Expression cells and click **Modify Expression**. Use the features of the dialog to compose and test the expression. When the expression works as desired, copy it to the clipboard, cancel the Link Wizard, then open the navigator item Properties dialog and paste the expression into the Item Symbol **Value** field.

After you enter the symbol name/value pair, the pair is evaluated. If the pair is valid, you are prompted to click **OK**. If the pair is not valid, a Syntax Error dialog is displayed, the title of the dialog indicates whether the name or value is in error, click **OK** and correct the name or value.

Examples

- Name: IMSID
  Value: IB1H

  This value resolves to IB1H.

- Name: IMSID.STA
  Value: TOKEN($NODE$,":",2) + "STA"

  This value resolves to *xxxx*STA, where *xxxx* is the second token from a node name value that is extracted from the execution context and having the format aaa:*xxxx*:bbb.

Assigned member options

Click ▦ **Options** in the **Assigned Members** area of the custom **Navigator item properties** to determine how the branch is populated.

All options are enabled by default.

☑ **Add branch-level items using managed system group names**

Add the branch items directly below the current item using the managed system group name as the branch item label.

In example 1, ▦ **Georgia group** is the managed system group name, which is absent when this option is disabled in example 2.

☑ **Include attribute-level items of managed systems**
Show the attribute-level child items of this Navigator item. These items are shared with the items found in the Navigator Physical view at the 🖥 Attribute level.

In example 2, 🖥 **Disk**, 🖥 **Memory**, and the other attribute level items are present. These items are absent when this option is disabled, as shown in example 3.

☑ **Include host-level items associated with managed systems**
Include directly below the current item the unique 🖳 System-level items (computer or node) that are associated with the managed systems being attached to the custom Navigator view.

In example 3, 🖳 **ATL-INTEL-X86-1** (Atlanta Intel X86 computer) is a node-level Navigator item, which is absent in example 4 when this option is disabled.

Examples:

1. All options selected
    ☑ Add branch-level items using managed system group names
    ☑ Include attribute-level items of managed systems
    ☑ Include host-level items associated with managed systems

    📁 **Shipping**
        📁 **Eastern Region**
            ▦ **Georgia group**
                🖳 **ATL-INTEL-X86-1**
                    ◉ **Windows OS**
                        🖥 **Disk**
                        🖥 **Memory**
                        🖥 ...

2. Managed system group names excluded
    ☐ Add branch-level items using managed system group names
    ☑ Include attribute-level items of managed systems
    ☑ Include host-level items associated with managed systems

Removes ▦ group-level items:

    📁 **Shipping**
        📁 **Eastern Region**
            🖳 **ATL-INTEL-X86-1**
                ◉ **Windows OS**
                    🖥 **Disk**
                    🖥 **Memory**
                    🖥 ...

3. Managed system group names and managed system child items excluded
    ☐ Add branch-level items using managed system group names
    ☐ Include attribute-level items of managed systems
    ☑ Include host-level items associated with managed systems

Removes ▦ group-level items, as well as the 🖥 child items of the ◉ Windows OS managed system:

    📁 **Shipping**
        📁 **Eastern Region**
            🖳 **ATL-INTEL-X86-1**
                ◉ **Windows OS**

4. No options selected
  - ☐ Add branch-level items using managed system group names
  - ☐ Include attribute-level items of managed systems
  - ☐ Include host-level items associated with managed systems

Removes ⊞ group-level items, as well as both the 🖳 child items and 🗗 host items of the ◉ Windows OS managed system:

📁 **Shipping**
  📁 **Eastern Region**
    ◉ **Windows OS**

Navigator List

After you right-click an item in the Navigator view or in the Navigator editor and click ⛓ **Show Navigator List**, this dialog box opens with a list of the Navigator views that the Navigator item belongs to.

The Navigator views shown include the current Navigator item. A number in parentheses after the view name indicates that there are multiple instances of this Navigator item in the Navigator view. For example, ⛓ Logical (2) means the Navigator view has two instances of this item.

To open another Navigator view, select it and click **Switch to**.

Chapter 10. Situations for event monitoring

Situations are definitions of conditions to test, such as slow transaction rates or invalid logon attempts. Use situations to raise alerts of certain conditions and to trigger single action commands or automated workflows.

Your user ID must have Modify Situation permission to open the Situation editor for creating and maintaining situations.

Overview

You can create monitoring specifications called situations to detect when specific conditions or events in your environment occur, thus raising an alert. Each situation is assigned to one or more managed systems that is to be monitored for specific conditions.

Predefined situations

Every monitoring agent comes with a set of predefined situations so that you can begin monitoring as soon as you begin using the Tivoli Enterprise Portal. Most of the predefined situations require you to enable them before they can monitor for events. This involves setting them to Run at startup if you want them turned on automatically with the agentYou can edit the thresholds of these situations to match your environment, as well as create your own situations.

Situations for pure events or sampled events

The event that opens for a situation can be *pure* or *sampled*. When the determination of the event is based on observations made at specific intervals, it is known as a sampled event. When the event is based on a spontaneous occurrence, it is known as a pure event. Therefore, situations for sampled events have an interval associated with them, while those for pure events have no interval. Another characteristic of sampled events is that the condition that caused the event can change, thus causing it to be no longer true. This is not the case with pure events: Pure events cannot change. Therefore, alerts raised for sampled events can transition from true to not true, while a pure event stays true when it occurs.

An example of a sampled situation is disk free space < 10%. The sample situation becomes true when the available disk space falls below 10% and later becomes false again after unneeded files are deleted and the disk space goes to 11% or higher. A situation that monitors for an invalid user logon attempt is a pure event: the event occurs when an invalid logon attempt is detected, and does not become a False event.

Situation editor

Open the Situation editor to see the definition of situations, edit them, or to create new situations. The Situation editor has tabs for:

- Composing the condition that will trigger the situation to fire, along with the interval between tests and the severity of the alert indicator (such as critical or harmless). When you open or create a situation, the column names that show in the formula editor are the names of the attributes being tested by the situation.
- Selecting the systems to test with the situation
- Writing expert advice to appear in the event results workspace

⚡ Entering a take action command to be run on the system when the situation fires

🖼 Specifying the severity and destination of events that are forwarded to an Event Integration Facility receiver. (Available only when the hub Tivoli Enterprise Monitoring Server has been configured to forward events to an event server. Refer to the event integration information in the *IBM Tivoli Monitoring Installation and Setup Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/ com.ibm.itm.doc_6.3fp2/install/itm_install.htm).)

🕐 Adding an Until modifier to control when an open event closes: after a period of time or when another situation is true

Situation editor opened from the toolbar (Ctrl+E) or from a Navigator item
You can open the Situation editor from the toolbar or from the Navigator item pop-up menu and there is a difference. The Situation editor is exactly the same for every aspect of the situation definition except for the fields for alert indicators in the Formula tab: Sound and State This behavior is intentional. The Situation editor opened with Ctrl+E shows all the situations in your environment, not filtered by application, and there is no context. Some administrators want to create many situations at one time and associate them with Navigators later, if at all. For example, you might want a situation running in the background that produces no visible alert on the Navigator but that might be embedded in another situation, might be used to activate a policy workflow, or to perform reflex automation. For an alert indicator to show in the Navigator, the Situation editor must be opened from the item where you want the alert.

Single-row and multiple-row attribute groups
You can build a situation with attributes from different groups belonging to the same monitoring agent so long as the attributes are from single-row attribute groups; one multi-row attribute group can be included.

Who can see situations
Situation definitions, both predefined and custom, are stored on the Tivoli Enterprise Monitoring Server. This means that any changes to situations made by any user are shared with all users on the same managed network if their user permissions include the ability to view situations and to access the managed application that is referenced in the situation.

Start situations manually or automatically
By default, situations are set to start monitoring as soon as the situation is saved and, subsequently, to start automatically when a monitoring agent that the situation is distributed to starts again after going offline. If you disable the **Run at startup** capability, you must start the situation manually.

Even if a situation is set to **Run at startup**, you can stop and start the situation manually with ▷ **Start Situation** and 🔴 **Stop Situation** if your user permissions include the ability to start and stop situations.

Even if a situation has been started manually, it can still be started automatically on other managed systems that connect to the same monitoring server. For example, if the managed system is added to the distribution after the situation was started manually, the situation starts running on the newly added managed system. This addition to the distribution might be caused by any of the following conditions:

• The monitoring agent came online after the situation was started.

- The managed system was added to a managed system group in the situation distribution.
- The managed system was added to a situation object group that the situation is a member of.

Command-line Interface

If you are using the Tivoli Monitoring command line **tacmd createSit** function (described in the *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/ com.ibm.itm.doc_6.3fp2/ic/landing_cmdref.htm)) for situation creation, you also can use the Situation editor to validate the specified attributes.

Private situations

Enterprise monitoring agents and system monitor agents can use locally defined situations to operate fully autonomously. These locally defined *private situations* are created in a private situation definition XML file. Private situations events come directly from the monitoring agent. You must place a private situation configuration file in the agent installation and restart the agent to enable this function. If you want to send an SNMP alert or EIF event when a private situation event is opened, then the SNMP trap configuration file or EIF event configuration file must also be in the agent installation.

Private situations on an enterprise monitoring agent have no interaction with or reporting of any kind to the monitoring server. Private situations and enterprise situations can run concurrently.

See "Private situations" in the *IBM Tivoli Monitoring Administrator's Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/ com.ibm.itm.doc_6.3fp2/adminuse/itm_admin.htm) for details.

Situation limitations

See the ITM Situation Limits technote.

Opening the Situation editor

There are several places from which you can open the Situation editor, with different capabilities related to the context. When the Situation editor is opened from the toolbar, for example, you can create situations for any agent type but you cannot associate them with Navigator items or set the event state or sound.

About this task

Open the Situation editor from one of these locations:

Procedure

- Right-click a Navigator item and select 🔲 **Situations**.

  The Situations tree shows the situations associated with the Navigator item. If you do not see the situation, click 🔧 **Situation Filter** to identify a broader range. The *state* and *sound* for the situation is also associated with the Navigator item. You can give the same situation a different state and sound for different Navigator item associations.

  You will not see this option at the Navigator Physical view's 📁 platform level (such as UNIX Systems).
- Click 🔲 **Situation Editor** in the toolbar.

  When opened through the toolbar, the Situation editor has no connection to the Navigator. You can see all the situations available in your monitored network,

you can edit or delete any of them, create new situations, and distribute them to managed systems. What you cannot do is associate situations with Navigator items or change their state or sound. This is also true when you open the Situation editor through the Manage Situations window.

- Right-click a situation in the event flyover list and select ☑ **Edit Situation**.

  The situation opens in the Situation editor so you can read the definition and edit it. You cannot associate (or disassociate) the situation with a Navigator item, nor can you change its state or sound here.

- Right-click an open event in the situation event console view and select ☑ **Edit Situation**.

  The situation opens in the Situation editor so you can read the definition and edit it. You cannot associate (or disassociate) the situation with a Navigator item, nor can you change its state or sound.

- In the Manage Situations at Managed System window (opened through the pop-up menu of either the ▦ Enterprise or ◙ agent Navigator item), select a situation row and click ☑ **Edit Situation**.

- While editing a policy workflow in the Grapher view, right-click a situation activity and click ☑ **Edit Situation**.

Filtering the situations tree

When you open the Situation editor from anywhere except the toolbar, only the situations associated with the current Navigator item are displayed in the Situation tree. You can filter the tree to show situations that are not associated or to see all the situations that have been defined for the monitored application.

Before you begin

You must open the Situation editor from a Navigator item for the ⮊ **Set Situation filter criteria** tool to be available.

About this task

The tree lists the situations that are associated with the Navigator item (or all situations in your monitoring environment if you opened the Situation editor from the toolbar). For example, if you opened the Situation editor from the Enterprise Navigator item, you would see all the situations that are distributed and associated with the Tivoli Enterprise Monitoring Server; if you right-clicked a system item, you would see all situations that were distributed to that managed system. If you open the Situation editor from a custom Navigator view, the situations are those associated with the managed systems that were assigned to the Navigator item.

Take these steps to control which situations appear in the tree:

Procedure

1. Click ⮊ **Situation Filter**.
2. Select any combination of these options:

Option	Description
Associated with Monitored Application	Show all situations that were written for this type of agent, regardless of where they are distributed or if they are distributed. Select this box if you want to see situations that have not been distributed to managed systems.
Eligible for Association	Show all situations that have been distributed to the managed system (or systems) assigned to this Navigator item and that have not been associated with it. If you are planning to associate situations with this Navigator item, select this box to see those situations that have not yet been associated.
Associated with this object	Show all situations associated with this Navigator item. If you are planning to associate situations with this Navigator item, clear this check box and select ☑ **Eligible for Association** to see only those situations that have not yet been associated.

Results

After you click **OK**, the tree shows the situations specified by the filter. If the results are not what you expected or need, return to the Show Situations window and try another setting.

Situation editor

Situations notify you when an event occurs on a managed system. Your monitoring product provides a set of predefined situations for your immediate use. You can view and edit these situations and create others.

The left frame of the Situation editor is a tree of the situations associated with the current Navigator item or the situations of all installed monitoring products if you opened the editor from the toolbar; the right frame shows this user assistance until you create or select a situation.

In the situation tree, click:

or Situation name to open the definition for viewing and editing. (indicates the situation is *not* distributed.)

Set Situation filter criteria to identify a broader range of situations to show in the tree.

Create new Situation to create a new standard situation or correlated situation.

Create another Situation to copy the selected situation, with a new name and description.

Delete Situation to permanently remove the selected situation.

After selecting or creating a situation, click:

*ƒ*x "Formula" on page 259 to see, add attributes to, and edit the formula.

"Distribution" on page 261 to see and assign the managed systems where the situation is running.

"Expert Advice" on page 262 to create links to URLs or to write instructions to be read in the event results workspace.

"Action" on page 263 to send a command or a universal message as an

automatic, reflexive response when the situation becomes true.

▣ "EIF" on page 265 to forward situation events to an Event Integration Facility receiver. (Available when the Tivoli Enterprise Monitoring Server has been configured to forward events.)

◷ "Until" on page 266 to specify that the event be closed after a certain period of time or when another situation is true.

In the situation tree pop-up menu, click:

▯ Create New situation to create a new standard situation or correlated situation.

▯ Create Another situation to copy the selected situation, with a new name and description.

▶ Start or ⏺ Stop the situation. Be careful when stopping a situation, because it stops on all managed systems where it is running.

▨ Delete the situation. The situation stops running on all managed systems where it was distributed, and then is deleted.

⇄ Associate the situation with the current Navigator item to enable an alert indicator of the selected state to be applied to the Navigator item when the situation is true; or ⣿ Dissociate the situation from the current Navigator item.

⊘ Enable or ⊖ Disable the situation overrides that have been created for this situation, or ⊠ Remove All Overrides.

⊞ Show the situation groups that this situation is a member of.

The Situation editor buttons enable you to **Apply** your changes by saving and starting the situation, and exit with (**OK**) or without (**Cancel**) saving your changes. If you plan to edit multiple situations while in the Situation editor, click **Apply** to save your changes before selecting the next situation to edit or creating a new one.

Note: Open the editor from a Navigator item if you want to edit or assign a state or sound to a situation, or if you want to associate a situation with a Navigator item.

Tip: If you are using the command line **tacmd createSit** function for situation creation, you can use the Situation editor to validate the specified attributes.

Situation tree

The Situation editor tree shows the situations that are associated with the current Navigator item; or it shows all situations in your monitoring environment if you opened the editor from the toolbar.

If you right-click the Enterprise item, for example, you see all the situations that are distributed to the monitoring server; if you right-clicked a system item, you see all situations that were distributed to that managed system. If you access the editor from a custom Navigator view, the products and situations shown are those associated with the managed systems you have assigned to the Navigator item through the Navigator properties window.

You can also show all the situations running on the operating platform (such as Windows or Linux) or specify a subset. To change what situations show in the tree, do the following:

1. Click ⋙ **Situation Filter**.
2. Select any combination of the following three options. Select all three to see every situation in your installation.

☐ **Associated with Monitored Application** to see all situations that were written for this type of agent, regardless of where they are distributed. Clear this box if you want to see situations that have not been distributed . When you return to the Situation tree, any undistributed situations have a ▤ gray disabled icon.

☐ **Eligible for Association** to show all situations that have been distributed to the managed system or systems assigned to this Navigator item and that have not been associated with it. If you are planning to associate situations with this Navigator item, select this box to see those situations that have not yet been associated.

☑ **Associated with this object** to show all situations associated with this Navigator item. If you are planning to associate situations with this Navigator item, clear this check box to see only those situations that have not yet been associated.

3. Click **OK** to see the result of the filtering you set. If the results are not what you expected or need, return to the Show Situations window and try another setting.

Formula

The **Formula** tab defines the condition that will be compared to the attribute values sampled by the agent. If the comparison is true, an event is opened. If you opened the editor from the Navigator, you can also assign a state and a sound to the situation.

Name   The situation name can contain any UTF-8 characters and can be up to 256 bytes long. The number of characters allowed varies depending on the characters used, because some UTF-8 characters require more than one byte.

If you enter a character other than letters, numbers, underscores, or spaces in the situation name, a unique short name is created and shown in the Situation editor status bar and in the results of the CLI tacmd viewSit. For example, a situation named My #1 situation might show as Z_____310A58ACAE95465A. This does no harm and cannot be changed. The only way to change the short name is to create another situation from this one, give it a name with no special characters other than the underscore, and delete the original situation.

Description

The description is displayed in the Manage Situations at Managed System window and when you select the situation for embedding in another situation. The description contain any UTF-8 characters and can be up to 64 bytes long. The number of characters allowed varies depending on the characters used, because some UTF-8 characters require more than one byte.

Formula

This area shows the thresholds to be tested by the situation. Click a cell under a column to activate the editor and enable the formula bar.

⚙ Show Formula opens a graphical view of the formula.

The active cell (and the formula bar) has three fields for building an expression ▯ == ▯ :

- **ν** **Function** shows the type of function chosen for the attribute, which is set to **Value of expression** by default unless you chose a situation condition (indicated by ⊡). Click the function to open and select from a list of possible functions.

- ⊟ Operator shows the relational operator. Click the operator button to open and select from a list of operators: **==** (equal to), **!=** (not equal to), **>** (greater than), **>=** (greater than or equal to), **<** (less than), and **<=** (less than or equal to). Note that some operators are not available for certain functions, such as ⊞ **Scan for string within a string**, which is limited to **==** and **!=**.

- ⊟ This field shows the value to be compared with the values arriving from the agent. Click inside the cell to add or change the value. Attributes with enumerated values have a ⊡ list button to click for the possible choices. Embedded situations show **True**. Correlated situations show **True** and **False** choices.

Note: The Situation editor performs syntax checking on each expression you enter. If you enter an improper value, such as abc for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

Situation Formula Capacity reports what percentage of the maximum formula size has been consumed. Typically, one condition or column heading uses 10%.

Add conditions opens a window for choosing an attribute to add to the situation. When you create a situation with a single-row attribute group, you can also add attributes from one multi-row attribute group or from any number of other single-row groups belonging to this monitoring agent. For complex expressions, such as when you need to repeat an attribute name, you can double-click the value text box to see and select an available attribute from the Select Condition window.

Advanced opens a window with two options: **Situation Persistence** to set the number of times the situation remains true before an event is opened; and **Display Item,** for multiple row attribute groups, to enable the situation to continue to test all rows in the data sampling and open events.

Keep multiple expressions in the same row if they must all be met (Boolean AND logic) and on separate rows if any of them can be met (Boolean OR logic) to set the situation to true. For example, if you want a situation to fire when either disk time OR disk space used is at least 90%, the formula would look like the one shown here.

	% Disk Time	% Used	Server Name
AND	>= 90		== 'myserver'
OR		>= 90	== 'myserver'

Add more rows by right-clicking a row number and selecting ⊡ **Insert**.

You can add an attribute twice for an AND comparison.

Sampling interval

The sampling interval is set to 15 minutes by default. You can change it to as seldom as once in 999 days or as often as 30 seconds, although more

frequent samplings consume more resources. Click inside the day, hour, minute or second cell and type the number or click the up and down arrows.

Note that if the situation monitors for pure events, this field is disabled; rather than sampling at specific intervals, the situation waits for a pure event.

Sound

(Not available when you open the Situation editor through the toolbar.)

☐ **Enable**, when selected, causes the sound file to play when an event is opened or reopened for this situation. Click **Play** to hear the sound. Click **Edit** to open the sound properties for editing.

State ⊗ ⊗ ⚠ ⚠ ✅ ℹ ⊘

(Not available when you open the Situation editor through the toolbar.)

This is the state indicated when the conditions have been met and the situation becomes true. An event indicator for the state, such as ⊗ Critical or ▪ Informational overlays the Navigator item icon. To specify a state, click the list button and select one.

☑ **Run at startup**

Starts monitoring as soon as the situation is saved with **Apply** or **OK**. If a monitoring agent to which the situation is distributed goes offline, the situation will start automatically when the agent starts again. Clear the box if you want to start the situation manually

Distribution

The Distribution tab is where you assign the managed systems where the situation should run. When you open the Situation editor from the Navigator item menu, the default distribution for a new situation is to the managed systems assigned to that item.

If the managed system has multiple agents of this type, then distribution will be to all of the agents. For example, a message queue situation written for a system with three message queues would be distributed to the three agents of this type on the system.

Assigned list

Shows the agents to which the situation is assigned.

Available Managed Systems

Shows the agents of the same type to which the situation can be distributed, such as all UNIX agents for a UNIX situation. To add a managed system to the distribution, click the name, then ◀ move it to the **Assigned** list.

Available Managed System Groups

Your product comes with one managed system group named for the product, such as *MS_SQL_SERVER for IBM Tivoli Monitoring for MS SQL Server managed systems, and you can have custom lists created through the Object group editor. To add a managed system group to the distribution, click the name, then ◀ move it to the Assigned list.

Note: The situation tree shows a ▬ disabled icon for a situation that has not been distributed.

Expert Advice

When a situation becomes true, an event indicator lights up on its associated Navigator item. Move the mouse pointer over the indicator, and click one of the event links to see the event results workspace, which includes **Expert Advice** for the event.

For the expert advice view, you can add comments or instructions for users, using the built-in HTML editor, or insert a link to an existing HTML page.

Adding expert advice in HTML text
> The text you enter in the expert advice view of the event results workspace is converted to an HTML document. You can format the text, such as to bold or italicize a word, using the proper HTML code. For example, **bold** and <i> *italic* </i> . And you can add links, using HTML tagging, as in <a href "http://www.ibm.com">IBM . Use `Preview` to check that your coding is correct. The limit is 512 characters. If you must go beyond that, consider linking to another HTML page.

Opening a web page
> If, instead of entering HTML text, you want to have a particular web page opened automatically in the expert advice view, simply enter the URL without any other text. For example, if you have only `http://www.ibm.com` for the expert advice text, the view will open to that website. Use **Preview** to test the URL.

Using conditional expressions
> The conditional expressions you can build in the Link Wizard customization page can also be included in expert advice using text scripting language. The syntax is very specific, so it is best to gather the attribute reference or compose the expression in the Link Wizard and copy it to the expert advice.
>
> The conditional expressions you can build in the Link Wizard customization page can also be included in expert advice using text scripting language. The syntax is very specific, so it is best to gather the attribute reference or compose the expression in the Link Wizard and copy it to the expert advice. **Preview** to check your coding is correct and does what you had intended. Example:

```
search = INPUT("What do you want to look for?", "enter");
  "<BODY BGCOLOR='#e6e6fa'>Searching database for <b>"+search+"</b>...
  <META HTTP-EQUIV='refresh' CONTENT='1;
  URL=http://searchenginesite.com/search?q="+search+"'>"
```

> The code creates the variable `search` from the user's response to, `What do you want to look for?` , then puts it in the URL for the search engine. The META `refresh` tag is used to switch to the website automatically after one second. The `INPUT` function is available in the Link Wizard. Example:

```
name = $EVENT:ATTRIBUTE.ISITSTSH.SITNAME$;
node = $NODE$;

"<BODY BGCOLOR='##FEDCEF'>
  <div align=center><h3><i>This is expert advice for </i>
  </h3><h2>"+name+"</h2> on "+node+ "<p>Please notify your system
  administrator that we have a <b>"+name+"</b> alert. <p>To learn more,
  click here: <a href=http://www.searchsite.com/search?q="+name+">
  IBM Tivoli Advice </a></div>"
```

The code in this example inserts the name of the situation and the name of the managed system in the text. It then provides a link to `searchsite.com` , which searches for the situation name and shows the results automatically. You can copy and paste this example into your expert advice HTML editor, replacing `searchsite.com` with the search-enabled web page containing your organization's run book procedures.

As well as the situation name given in this example, you can create a variable for any other event attribute:

Attribute	Variable
situation name	$EVENT:ATTRIBUTE.ISITSTSH.SITNAME$
monitoring server name	$EVENT:ATTRIBUTE.ISITSTSH.NODE$
managed system name	$EVENT:ATTRIBUTE.ISITSTSH.ORIGINNODE$
display item (if set)	$EVENT:ATTRIBUTE.ISITSTSH.ATOMIZE$
global timestamp	$EVENT:ATTRIBUTE.ISITSTSH.GBLTMSTMP$
local timestamp	$EVENT:ATTRIBUTE.ISITSTSH.LCLTMSTMP$
status (Y/N)	$EVENT:ATTRIBUTE.ISITSTSH.DELTASTAT$

Action

Use the **Action** tab to send a command to a managed system or a message to the universal message console view when the situation becomes true. You might want to log information, trigger an audible beep, or stop a job that is overusing resources.

The command can be a single action or a script containing several commands. No feedback is received after the command or user notification is sent.

◉ System Command
○ Universal Message

Note: For an action to run, the situation must not be used in a policy. If the Action fields are disabled, it means your monitoring agent does not support this feature.

Tip: You can use Tivoli Enterprise Monitoring Web Services to send Simple Object Access Protocol (SOAP) requests between hub monitoring servers. See the *IBM Tivoli Monitoring Administrator's Guide* for details.

System command

When you select **System Command**, this field is displayed for you to type a command to issue at the system. The command you enter depends on the application and operating platform you are running. For example, you can enter a UNIX shell command to run at a UNIX system.

Example: Windows has a command for sending a message to a system: `net send` *computer_name* `message`, as in `net send MYPC Disk is 85% full. Delete some files.`

Example: UNIX has the mail command for directing messages: `mail node userID text`.

Note: To use the ":" colon character in the command, precede it with a "\" backslash. Otherwise, the command will truncate at the ":" colon. For example, to specify `c:\something` , you would enter `c\:\something` .

The **Attribute Substitution** button opens a window for you to choose an attribute. The attribute name you select is replaced by the actual value during execution of the action command. Note that attribute substitution is not possible for correlated situations. Example:

```
echo System: &NT_Logical_Disk.Server_Name, Disk: &NT_Logical_Disk.Disk_Name,
  Space Available: &NT_Logical_Disk.%_Free%|mail userID .
```

This command sends an email that reads, "System: MyServer, Disk: D:, Space Available: 15

Universal message

You can generate a universal message that can be seen in a universal message console view.

Category is the category the message must fall under, such as "Non-essential" or "SiteCritical". Type a one-word term of up to 16 characters to categorize the message.

Severity is the importance of the message, such as "High" or "1". Type a one-word term of up to 8 characters that conveys the severity of the message.

Message is the text to display when the situation occurs. Type a message of up to 245 characters. You can include an attribute value in the message by selecting **Attribute Substitution**.

Attribute Substitution opens a window for choosing an attribute. The attribute name is replaced by the actual value during execution of the action command.

If the condition is true for more than one monitored item:

- If the sampling from the monitoring server returns multiple rows of data and more than one row satisfies the condition, you can choose to issue the command on only the first row that meets the condition or once for each row that meets the condition. For example, if the situation fires when any process uses over 80% of the CPU, you can issue a command to terminate just the first process that meets this criteria or all processes.

    ◉ **Only take action on first item**
    ○ **Take action on each item**

Note: If the situation has a Display Item set and the "Only take action on first item" option is selected, the action is taken once for each unique display item in the rows of data. For example, if there were 100 total rows that qualified and there were 6 unique display items, the situation would fire 6 times. See "Formula" on page 259 for more information on display items.

Where should the Action be executed:

- Determines where to perform the action: on the system where the agent resides or at the Tivoli Enterprise Monitoring Server to which it is connected.

    ◉ **Execute the Action at the Managed System (Agent)**
    ○ **Execute the Action at the Managing System (TEMS)**
    If you are sending a message to the universal message console view, select this option.

If the condition stays true over multiple intervals:

- Determines when to take action on a true situation. If you want the command to run only once, and not every time incoming data matches the condition, then

leave it at the default setting ◉ **Don't take action twice in a row (wait until situation goes false then true again)**. If the situation remains true over successive intervals, the command runs only the first time the situation is true and not again until it becomes true from being false in the previous interval.

- To issue the command on every true situation evaluation, even if no false states intervene, select ○ **Take action in each interval** Enable this option to invoke the command when the situation is true in an interval, regardless of its state in the previous interval.

Suppose your UNIX system has recently been plagued by insufficient memory. You are not certain your system really needs more memory resources; maybe a poorly designed application or process is causing the problem. You want to gather system data to analyze the cause of the memory problem, but you do not want to allocate personnel to gather the data. You can write a situation to monitor excessive page-ins and page faults and, when either occurs (indicating a memory bottleneck), write the process name and ID to a log for analysis by an application at a later time.

EIF

Use the EIF (Event Integration Facility) tab to forward situation events to one or more EIF receivers (such as a IBM Tivoli Enterprise Console event server or Netcool/OMNIbus Probe for Tivoli EIF).

This tab is available only if the hub monitoring server has been configured to forward events. See the *IBM Tivoli Monitoring Installation and Setup Guide* for details.

☐ **Forward Events to an EIF Receiver**
> Select this check box to specify that situation events should be forwarded to one or more EIF receivers. If this option is selected, an EIF event is sent for each situation event.

EIF Severity
> Select the severity for forwarded situation events. <Default EIF Severity> causes the event severity that is set for the situation at this Navigator item to be applied to the EIF event when ☑ **Forward Events to an EIF Receiver** is selected. Default EIF severity mapping:

Tivoli Monitoring situation event severity	Tivoli Enterprise Console Server event severity	Tivoli Netcool/OMNIbus ObjectServer event severity
⊗ Fatal	not used	60 = Critical
⊗ Critical	⊗ Critical	50 = Critical
⚠ Minor	not used	40 = Minor
⚠ Warning	⚠ Warning	30 = Warning
☑ Harmless	☑ Harmless	0 = Clear
ℹ Informational	☑ Harmless	not used
◈ Unknown	ℹ Informational	10 = Indeterminate

EIF Receivers
> Specify the EIF destinations for the forwarded events.
>
> The **Assigned EIF Receivers** list shows the destinations to which forwarded events will be sent, and the **Available EIF Receivers** list shows all of the defined EIF destinations. (EIF destinations are defined in the

Manage Tivoli Enterprise Monitoring Services application, described in the *IBM Tivoli Monitoring Installation and Setup Guide*, or with the **tacmd createEventDest** command, described in the *IBM Tivoli Monitoring Command Reference*.)

Each destination represents one or more EIF receivers. If you move the mouse pointer over an entry in the list, the hover help shows the server and port information specified for the destination.

To add a destination, select it from the **Available EIF Receivers** list and then ◄ move it to the Assigned list. (After selecting the first destination, you can use Ctrl+click to select other destinations or Shift+click to select all destinations between the first selection and this one.)

To remove destinations from the list, select them from the **Assigned EIF Receivers** list and ▶ move to the Available list.

If you want to customize how situation events are mapped to EIF events, click **EIF Slot Customization**

Note: If you are forwarding events, when an active situation is deleted, the Stop and Delete events that are sent to the event server use the situation short name and not the long name. This is also true for any messages that are sent to the message log view.

Until

The *Until* clause provides a way to automatically close an event by turning a true situation to false when another situation is true or when the close interval expires. This is particularly useful for pure events that might occur frequently.

Close this Event when either occurs
   □ **Interval Expires** When this option is cleared, the event does not close until the situation becomes false. Select this option to close the event automatically after the specified interval, 15 minutes by default. To change the interval, click inside the days, hours, minutes or seconds cell, then ▾ adjust up or down, or type the number.

   This option must be used only for Pure Events in cases where processing is occurring outside the portal client, such as situation action commands or by an event processor like NetCool/OMNIbus. It should never be used for sampled situations. The underlying condition will continue to exist even though there is no indication. The event will open again only after the condition become false and then goes back to true. For sampled situations, there are three ways to close an event:
   • Resolve the issue
   • Stop the situation
   • Change the situation distribution

   □ **Another Situation is TRUE**, when selected, closes the event for this situation when the selected situation is also true. The situations available are for the same monitored application that this situation is written for or that use the common attributes. The two events must have the same managed system names and Display Item if present in either situation. For sampled situations, after an event is closed based on the Until setting, no further events are opened even if the situation continues to evaluate to true. The situation must evaluate to false, then true again before another event is opened. However, if the situation also has persistence set (click **Advanced** in the Formula tab), and continues to be true after the event has

been closed, another event is opened when the persistence count has passed.

This option is designed for sampled events. It can work for pure events if they arrive slowly. Events arriving in sudden bursts can be reordered by the monitoring server. In some agents a burst of activity can lead to loss of events. Either of those two cases will result in processing failures. As a result any such usage must be tested extremely carefully.

Predefined situations

Tivoli Monitoring products have situations that are ready to use as soon as you start the portal client. Also provided are situations that can be applied to any managed system and situations for the monitoring server.

Some predefined situations are already distributed and set to start automatically when the monitoring server is started. And some of those situations might not be associated with a Navigator item, which means an event is opened when the situation becomes true and an entry is written in the message log, but no event indicator appears in the Navigator. Other predefined situations might not be distributed or might not be set to run at startup.

When setting up event monitoring, review the predefined situations and determine which ones you would like to have run automatically and whether any threshold values need to be changed. If the situation is not associated with a Navigator item and you want an event indicator to display there when an event is opened, set the alert state and, optionally, the sound to play.

All Managed Systems situations

No matter where you open the Situation editor from, the All Managed Systems situations are available. Except for the MS_Offline situation, which uses the Managed Systems attributes, these predefined situations use the Local Time attributes.

The default setting for these situations is to not **Run at startup**. Distribute these situations to the managed systems you want to monitor.

Historical_Exports_Critical

The upload interval from the managed system to the warehouse proxy is critical. When the queue reaches 100, the historical upload will take longer than its configured upload interval. The situation as it appears in 🦋 Show Formula when Show detailed formula is selected:

(#'KRAHIST.EXPTINUSED'>=95)

Historical_Exports_Failure

The historical collection export failed. This is the Warehouse Proxy agent return code. A return code that is not a zero indicates failure. The situation as it appears in 🦋 Show Formula when Show detailed formula is selected:

(#'KRAHIST.LSTEXPSTAT'!=0)

Historical_Exports_Warning

The upload interval from the managed system to the warehouse proxy is approaching critical levels. When the queue reaches 100, the historical upload will take longer than its configured upload interval. The situation as it appears in 🦋 Show Formula when Show detailed formula is selected:

(#'KRAHIST.EXPTINUSED'>=70 AND #'KRAHIST.EXPTINUSED'<=95)

Historical_RowsCorrupted

Errors were found in the short term history files. The corruption has been corrected, but the corrupted rows could not be exported to the warehouse proxy. The situation as it appears in ⚙ Show Formula when Show detailed formula is selected:

```
(#'KRAHIST.ROWSERRORS' > 0)
```

MS_Offline

The managed system offline situation opens an event if one of the managed systems gets disconnected or goes offline for some reason. It runs at one-minute intervals. When this situation is associated with the Enterprise item, any managed system in the network that goes offline will cause an event to open. The situation as it appears in ⚙ Show Formula:

```
( Status == *OFFLINE AND Reason != 'FA')
```

If an event is opened for this situation, it means the agent is not connected to its monitoring server, has stopped running, or is in error. If all the agents connected to the same monitoring server are offline, it is likely that the server is offline and not the agent.

NonPrimeShift

NonPrimeShift checks for times after 5 PM and before 8 AM from Mondays to Fridays. This situation is useful for embedding in a situation or policy that you want an event opened for only outside the prime shift hours. The situation as it appears in ⚙ Show Formula:

```
(Day Of Week >= Monday AND Day Of Week <= Friday AND Hours < 08) OR
(Day Of Week >= Monday AND Day Of Week <= Friday AND Hours >= 17)
```

PrimeShift

This situation is useful for embedding in a situation or policy that you want an event opened for only outside the prime shift hours. The situation as it appears in ⚙ Show Formula:

```
( Status == *OFFLINE AND Reason != 'FA')
```

QOMEGAMON_ONLINE

The situation as it appears in ⚙ Show Formula:

```
( Status == *OFFLINE AND Reason != 'FA')
```

Weekday

This situation is useful for embedding in a situation or policy that you want an event opened for only outside the prime shift hours. The situation as it appears in ⚙ Show Formula:

```
( Status == *OFFLINE AND Reason != 'FA')
```

Weekend

This situation is useful for embedding in a situation or policy that you want an event opened for only outside the prime shift hours. The situation as it appears in ⚙ Show Formula:

```
( Status == *OFFLINE AND Reason != 'FA')
```

Tivoli Enterprise Monitoring Server situations

The Enterprise Navigator item has a set of predefined situations. They appear within the Tivoli Enterprise Monitoring Server entry of the Situation editor.

Deploy_Failed

This situation becomes true and an event is opened when one or more agent deployments has failed. Default: Do not ☐ **Run at startup**. ⚙ Show Formula:

```
( Status == Failed)
```

Deploy_Retrying

This situation becomes true and an event is opened when one or more agent deployments is in the FailedRetry state, which means it has failed an attempt to deploy and is making another attempt. Default: Do not ☐ **Run at startup**. 🔍 Show Formula:

```
( Status == 'Failed Retrying')
```

Hub_Catalog_Critical

This situation becomes true and an event is opened when an application catalog that is installed at a remote monitoring server is *not* installed at the hub monitoring server. This can happen when application support has not been selected for the hub during installation. Default: Do not ☐ **Run at startup**. 🔍 Show Formula:

```
( Status == Missing At Hub) AND (Catalog != 'PDSSTATS')
```

Hub_Catalog_Warning

This situation becomes true and an event is opened when an application catalog that is installed at a remote monitoring server is newer than what is installed at the hub monitoring server. This is a warning that the application support for the component is out of date. This might not be a problem if the hub has not been upgraded yet. Default: Do not ☐ **Run at startup**. 🔍 Show Formula:

```
( Status == RTEMS Newer)
```

Remote_Catalog_Warning

This situation becomes true and an event is opened when an application catalog that is installed at the hub monitoring server is not also installed at a remote monitoring server; or an application catalog that is installed at the hub is newer than what is installed at a remote monitoring server.

This might not be a problem if the remote monitoring server has not been upgraded yet. Also, the remote system might be intentionally not supporting the subject application, or might only be supporting an older version of the application. Default: Do not ☐ **Run at startup**. 🔍 Show Formula:

```
((Status == Missing At Hub) AND (Catalog != 'PDSSTATS')) OR
(Status == Hub_Newer)
```

TEMS_Busy

This situation becomes true and an event is opened when the Tivoli Enterprise Monitoring Server is continually busy. Default: Do not ☐ **Run at startup**. 🔍 Show Formula:

```
( Product == 'EM' AND Reason == 'H')
```

Creating a situation

You can create a situation to alert you of certain conditions and to trigger single action commands or complex automated workflows.

About this task

Take these steps to create a situation:

Procedure

1. If the Situation editor is not open, open it using one of the methods described in "Opening the Situation editor" on page 255.

2. In the Situation editor, click ⬚ **Create new Situation**.
3. Type a name for the situation in the **Name** field and, optionally, a description in the **Description** field. The name and description can contain any UTF-8 characters: up to 256 bytes for the name and up to 64 bytes for the description. The maximum varies depending on the characters used because some UTF-8 characters require more than one byte. If you use any special characters such as pound sign (#) or ampersand (&) in the name, a unique short name is generated and displayed in the Situation editor status bar and in the results of the CLI tacmd viewSit.
4. If this will be a standard situation, select the type of monitoring agent from the **Monitored Application** list and click **OK**.
5. If this will be a correlated situation, select `Correlated Situation` from the **Type** list, click **OK**, and continue to step 5 of "Correlating situations" on page 314.

Defining the condition

Define the condition to test for in the Situation editor Formula tab.

Before you begin

This procedure assumes that you are creating a situation ("Creating a situation" on page 269) and the Select Condition dialog box is open. If you are editing a situation, this dialog box is opened after you click **Add Conditions** in the Formula tab.

About this task

Take these steps to compose the situation formula.

Procedure

1. In the Select Condition dialog box, select the type of condition that you want to define:
   - **Attribute Comparison**, then select an attribute group and one or more attributes. Any attributes that are inappropriate for use in situations do not appear in the list. For example, you might see an attribute in a table view that you cannot see when creating a situation with the same attribute group.
   - **Situation Comparison**, then select the situation to embed.
   - Use Ctrl+click to select additional items from the list or Shift+click to select all items between and including the first selection and this one. Attributes or situations that have already been chosen are colored blue. You can add additional attributes or situations later using the **Add conditions...** button on the 𝑓𝑥 **Formula** tab.

   After you select the attributes or situations to add to the formula and click **OK**, they appear as column headings in the formula editor.
2. For each attribute selected, click inside the cell under the column heading and compose an expression ▣ == ▢ , which consists of a function, a relational operator, and the value to test for:
   - To change the function, click ▣ Value (or ▣ Compare date/time for a time attribute) and select one from the list. (See Appendix A, "Formula functions," on page 415 for a description.)
   - To change the relational operator, click ▣ equal and select one from the list.

- Click inside the ☐ text box and enter or select the value.

  Keep multiple expressions in the same row if they must all be met (Boolean AND logic) and on separate rows if any of them can be met (Boolean OR logic) to make the situation true.

3. To add another attribute to the situation or to embed another situation, click **Add Conditions** and select the condition type and attributes or situations to add.

4. To add situation persistence or a display item to the situation, click **Advanced**:
   - **Situation Persistence** Specify how long a true situation must persist (the number of consecutive samples evaluating to true) before an event is opened. Persistence is useful for determining whether a condition is solid or intermittent. You can use it with the take action feature to start a command if the condition persists for the designated time (*persistence* x *interval*). You can also use this feature to control when an event indicator appears on the associated Navigator item or, in a situation that is referenced in a policy, to control workflow.
   - **Display Item** Normally, when a situation encounters a value in a row that meets the condition, an event is opened and remains open across sampling intervals until the situation is no longer true. No more events for the situation on this managed system can be opened until the event has been closed. However, if you set a display item for a multiple-row attribute group, the situation continues to look at the other rows in the sampling and opens more events if other rows qualify. This optional addition to a situation also enables you to quickly distinguish events in the message log and event console views. Only certain attributes can be designated as display items. Any valid display items are displayed in the list.

5. To specify how often the condition should be tested, accept the default of 15 minutes in the **Sampling interval** field or adjust the value. This field is disabled if the situation monitors for pure events. (Pure event occurs automatically, such as when a paper out condition occurs on the printer or when a new log entry is written. Situations written to notify you of pure events wait until an event occurs rather than sampling at specific intervals and remain true until they are manually closed or automatically closed by an UNTIL clause.)
   - Click inside the day, hour, minute, or second cell and type the number or click the up and down arrows.
   - You can change it to as seldom as once in 999 days or as often as 30 seconds, although more frequent samplings consume more resources.

6. To start the situation manually rather than automatically, clear the ☑ **Run at startup** check box. By default, situations are set to start monitoring as soon as the situation is saved and, subsequently, to start automatically when a monitoring agent that the situation is distributed to starts again after going offline. See Start situations manually or automatically for more information.

7. If you opened the Situation editor from a Navigator item, you can also assign a severity indicator or a sound to play when an event opens:
   - Assign a different severity by selecting an alert indicator from the **State** list.
   - Select a sound to play by checking ☑ **Enable** *filename* in the **Sound** area. Click **Edit** to change the sound file or its properties. Click **Play** to listen to the sound.

8. Select other tabs in the Situation editor as needed to complete the situation definition:

- Distribute the situation to the managed systems on which you want it to run. See "Distributing the situation" on page 273.
- Add expert advice to the situation. See "Writing expert advice" on page 275.
- Specify an action to take or a message to send to the Universal Message Console when the situation is true. See "Specifying an action" on page 277
- Forward events that open for this situation to an Event Integration Facility receiver such as the Netcool/OMNIbus Probe for Tivoli EIF. See "Forwarding the event to an EIF receiver" on page 280.
- Specify how long a situation shall remain true. See "Adding an Until modifier" on page 279.

9. When you are finished composing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor.

   The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

   If you opened the Situation editor from the toolbar, the situation you just created is not associated with any Navigator items and no event indicator is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

   Situations are stored at the hub monitoring server. In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error occurs to ensure that one situation does not overwrite another.

10. If the Distribution Change window is displayed, click **OK** to add the managed systems shown in the **Required Managed Systems** list to the distribution for this situation.

Playing sounds when an event is opened

Every situation has a state assigned to its associated Navigator item. You can also assign a sound file that plays whenever the situation becomes true for this Navigator item, and you can use different sounds for different Navigator items associated with the same situation.

About this task

Take the following steps to assign a sound to a situation.

Procedure

1. Right-click the Navigator item to which the situation is associated and select ☒ **Situations**. The Situation editor opens with the list of situations that have been associated with this Navigator item.

2. Select the situation to which you want to assign a sound. If you do not see the situation in the list, click ☰ **Situation Filter**, then select ☑ **Eligible for Association**, and click **OK**.

   The right frame displays the **Formula** tab for the selected situation. In the **Sound** area, the name of the default sound file for the state assigned to the Navigator item.

3. To use the default sound, select ☑ **Enable** The default sound for the state assigned to the Navigator item will play when this situation becomes true and opens or reopens an event.

4. Click **Play** to hear the currently selected sound; click **Edit...** to select another sound file or to edit the sound properties. Clicking **Edit...** opens the Sound window. The text field in the File area shows the name of the currently selected file (the default file for the state assigned to the Navigator item if no other has been selected).

   - Click **Browse** to select an alternative file. Default sounds are stored in the `sound` or `user` subdirectory on the system on which the Tivoli Enterprise Portal is installed.

   - Apply any of the following options to the selected file:

     Protect from Global Default overrides
     If selected, this option keeps the sound file you have assigned to this situation, even if you have selected a global option. If this option is not selected, any applicable global sound option you have set overrides the sound file shown here.

     Enable Sound
     If selected, this option use the sound shown in the File field. If this option is not selected, no sound plays when the situation becomes true. (You can also enable the selected file in the **Sound** area of the **Formula** tab.)

     Global Options
     This button opens the Sound - Global Options window where you can specify options for applying the sound to all situations of the same state or type as this one.

   - Click **OK** to save the selections and return to the **Formula** tab.
5. Click **Apply** to save your changes without closing the Situation editor, or click **OK** to save the changes and close the editor.

Customizing sounds with your audio files

You can add your own WAVeform audio format files to your Tivoli Enterprise Portal installation to customize the sound for a situation or situation type. Waveform audio is a format for storing sound in files, developed jointly by Microsoft and IBM.

About this task

On the computer where the Tivoli Enterprise Portal Server is installed, copy your custom WAV files to the `user` subdirectory of the `sound` directory:

```
<itm_installdir>
 CNB (cw on Linux or UNIX)
      📁 classes
      📁 candle
           📁 fw
                📁 resources
                     📂 sounds
                          📂 user
```

Use the **Edit...** button in the **Sound** area and then the **Browse** button on the *situation_name* - Sound window. to select your own WAV file.

Distributing the situation

The distribution of a situation determines where it will run.

About this task

The situation tree displays ![icon] when a situation is distributed, and ![icon] when a situation is *not* distributed.

Note: If you are running multiple versions of a monitoring agent, you will be able to distribute a situation only to the managed systems that the version supports.

Take the following steps to distribute a situation:

1. In the Situation editor, select the ![icon] **Distribution** tab.

   If it is not already, open the Situation editor and select the situation whose distribution you want to specify.

2. To add to the **Assigned** list, select from the **Available Managed Systems** or **Available Managed System Groups**, or both, and click ‹ .

   Available Managed Systems shows the monitoring agents to which the situation can be distributed.

   Available Managed System Groups Each monitoring agent product comes with one managed system group named for the product, such as *MS_SQL_SERVER for IBM Tivoli Monitoring for MS SQL Server managed systems, and you can have custom lists that were created in the ![icon] **Object group editor**.

3. If you want to override the situation formula for any of the assigned managed systems, select one or more managed systems or managed system groups from the **Assigned** list, and then click **Override Formula**.

   Note: Not all situations can be overridden. If you are editing a situation that cannot be overridden, the **Override Formula** button is not available.

4.
   Select other tabs in the Situation editor to:

   - Add expert advice to the situation. See "Writing expert advice" on page 275.
   - Specify an action to take or a message to send to the Universal Message Console when the situation is true. See "Specifying an action" on page 277
   - Forward events that open for this situation to an Event Integration Facility receiver such as the Netcool/OMNIbus Probe for Tivoli EIF. See "Forwarding the event to an EIF receiver" on page 280.
   - Specify how long a situation shall remain true. See "Adding an Until modifier" on page 279.

5. When you are finished composing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor.

   The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

   If you opened the Situation editor from the toolbar, the situation you just created will not be associated with any Navigator items and no event indicator is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

   Situations are stored at the hub monitoring server. In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error will occur, ensuring that one situation does not overwrite another.

6. If the Distribution Change window is displayed, click **OK** to add the managed systems shown in the **Required Managed Systems** list to the distribution for this situation.

Writing expert advice

Use the **Expert Advice** tab to add comments or instructions, up to 500 characters, for the user to see in the event results workspace. Or create a link to an existing page on the web which contains relevant information.

About this task

1. In the Situation editor, select the ⓘ **Expert Advice** tab.

   If it is not already open, open the Situation editor and select the situation to write expert advice for.

2. In the HTML editor, do one of the following :

   - Enter comments or instructions.

     The text is converted to an HTML document to be viewed from the event results workspace. You can format the text, such as to bold or italicize a word, using the correct HTML code. For example,

     `<b>`**bold**`</b> and <i>`*italics*`</i>`

     You can also add links, using HTML tagging, as in

     `<a href "http://www.ibm.com">IBM</a>`

     Click **Preview** to check that your coding is correct.

     You can also include conditional expressions in the Expert Advice. The syntax is very specific, so it is best to gather the attribute reference or to compose the expression in the Link Wizard and copy it to the expert advice. An expression can be a combination of literal values, references to values available from the source context, operators, and functions applied to values. Click **Preview** to check that your coding does what is intended.

   - Open a web page.

     If you would like to automatically open a web page in the Expert Advice view, enter just the URL without any other text in the HTML editor. For example, if the only entry in the editor is `http://www.ibm.com`, the view opens to that website. Click **Preview** to test the URL.

3. Select other tabs in the Situation editor to:

   - Distribute the situation to the managed systems on which you want it to run. See "Distributing the situation" on page 273.
   - Specify an action to take or a message to send to the Universal Message Console when the situation is true. See "Specifying an action" on page 277
   - Forward events that open for this situation to an Event Integration Facility receiver such as the Netcool/OMNIbus Probe for Tivoli EIF. See "Forwarding the event to an EIF receiver" on page 280.
   - Specify how long a situation shall remain true. See "Adding an Until modifier" on page 279.

4. When you are finished composing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor.

   The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

If you opened the Situation editor from the toolbar, the situation you just created will not be associated with any Navigator items and no event indicator is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

Situations are stored at the hub monitoring server. In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error will occur, ensuring that one situation does not overwrite another.

5. If the Distribution Change window is displayed, click **OK** to add the managed systems shown in the **Required Managed Systems** list to the distribution for this situation.

Example

An expression can be a combination of literal values, references to values available from the source context, operators, and functions applied to values. Click **Preview** to check that your coding does what is intended.

Search variable
> The code creates the variable "search" from what the user enters when prompted, "What do you want to look for?", and enters that value as part of the URL for the searchenginesite.com search engine. The META 'refresh' tag is used to switch to the website automatically after one second. The INPUT function is available in the Link Wizard.

```
search = INPUT("What do you want to look for?", "enter");"
  <BODY BGCOLOR='#e6e6fa'>Searching database for <b>"+search+"</b>...
  <META HTTP-EQUIV='refresh' CONTENT='1; URL=http://searchenginesite.com
  /search?q="+search+"'>"
```

Search variable with situation information included
> The code in this example inserts the name of the situation and the name of the managed system in the text. It then provides a link to *searchenginesite.com*, which searches for the situation name and displays the results automatically.

```
name = $EVENT:ATTRIBUTE.ISITSTSH.SITNAME$;
  node = $NODE$;
  "<BODY BGCOLOR='#FFFF80'><div align=center>
  <h3><i>This is expert advice for </i> </h3>
  <h2>"+name+"</h2> on "+node+ "<p>Please notify your system
  administrator that <b>"+name+"</b> fired. <p>To learn more, click here:
  <a href=http://www.searchenginesite.com/search?q="+name+">
  Tivoli Advice </a>"
```

> As well as the situation name given in the previous example above, you can reference any other event attribute in the code:

> **situation name**
>> $EVENT:ATTRIBUTE.ISITSTSH.SITNAME$

> **monitoring server name**
>> $EVENT:ATTRIBUTE.ISITSTSH.NODE$

> **managed system name**
>> $EVENT:ATTRIBUTE.ISITSTSH.ORIGINNODE$

> **display item**
>> $EVENT:ATTRIBUTE.ISITSTSH.ATOMIZE$

> **global timestamp**
>> $EVENT:ATTRIBUTE.ISITSTSH.GBLTMSTMP$

local timestamp	
	$EVENT:ATTRIBUTE.ISITSTSH.LCLTMSTMP$
status	$EVENT:ATTRIBUTE.ISITSTSH.DELTASTAT$

Specifying an action

Use the Situation editor **Action** tab to send a command to a system or a message to the universal message console when the situation becomes true.

You might want to log information, trigger an audible beep, or stop a job that is overusing resources.

About this task

A take action command can be a single action or a script of many commands. The Tivoli Enterprise Portal Server receives no feedback after it sends the command or notifies the user. To issue a command, the user ID must be authorized on the relevant system for the requested command. For example, to issue a TSO command, your user ID must be a valid TSO ID. The user ID must also be defined as a valid logon ID for the Tivoli Enterprise Portal Server. If an application prefix is used (such as MQ:), then the portal server logon ID will be used. However, for system level commands, the authorization of the monitoring agent will be used.

Note: The action will not be invoked if the situation is used in a policy.

Procedure

1. In the Situation editor, click the **Action** tab. If it is not already, open the situation editor and select the situation to which you want to specify reflex automation.
2. Select the type of action you want to take: System Command or Universal Message.
3. Specify the command or message.
   * For a system command, write the command exactly as you would enter it manually for the application and operating platform on which it is running.
   * For a universal message, write the category, severity and the text to display in the Universal message console view.
     - Category of messages to see (for example, critical, warning, information), up to 16 characters.
     - Severity of the message, one word of up to 8 characters, such as **1** or **high**.
     - Message text when the situation occurs, up to 256 characters.

   Attribute Substitution opens a window for choosing an attribute. The attribute name is replaced by the actual value during execution of the action command. Example:

   ```
   echo System: &NT_Logical_Disk.Server_Name, Disk: &NT_Logical_Disk.Disk_Name,
     Space Available: &NT_Logical_Disk.%_Free%|mail userID.
   ```

   which sends an email to the user ID indicated that reads, "System: MyServer, Disk: D:, Space Available: 15%."
4. If the condition is true for more than one monitored item, indicate when the action should be taken:
   * Select **Only take action on first item** to issue the command on only the first row that meets the condition.

- Select **Take action on each item** to issue the command once for each row of data returned that meets the condition.

  For example, if the situation fires when any process uses over 80% of the CPU, you can issue a command to terminate just the first process that meets this criteria, or you can issue a command to terminate all processes that meet the criteria.

  Note: If the situation has a Display Item set and the "Only take action on first item" option is selected, the action is taken once for each unique display item in the rows of data. For example, if there were 100 total rows that qualified and there were 6 unique display items, the situation would fire 6 times. See "Formula" on page 259 for more information on display items.

5. Specify where the action should be executed (performed):
   - Select **Execute the Action at the Managed System** to perform the action on the system where the monitoring agent resides.
   - Select **Execute the Action at the Managing System (TEMS)** to perform the action on the monitoring server to which the agent is connected

   If you are sending a Universal Message, the action takes place at the monitoring server.

6. Specify what to do if the condition stays true over multiple intervals:
   - Select **Don't take action twice in a row** to execute the command or universal message once, and not every time incoming data matches the condition.
   - Select **Take action in each interval** to invoke the command or issue the universal message when the situation is true in an interval, regardless of its state in the previous interval.

7. Select other tabs in the Situation editor to:
   - Distribute the situation to the managed systems on which you want it to run. See "Distributing the situation" on page 273.
   - Add expert advice to the situation. See "Writing expert advice" on page 275.
   - Forward events that open for this situation to an Event Integration Facility receiver such as the Netcool/OMNIbus Probe for Tivoli EIF. See "Forwarding the event to an EIF receiver" on page 280.
   - Specify how long a situation shall remain true. See "Adding an Until modifier" on page 279.

8. When you are finished composing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor. The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

   If you opened the Situation editor from the toolbar, the situation you just created will not be associated with any Navigator items and no event indicator is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

   Situations are stored at the hub monitoring server. In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error will occur, ensuring that one situation does not overwrite another.

9. If the Distribution Change window is displayed, click **OK** to add the managed systems shown in the **Required Managed Systems** list to the distribution for this situation.

Adding an Until modifier

Add an Until modifier to control when a true situation closes: after a period of time or when another situation is true or, if you select both, whichever occurs first.

Before you begin

When the Until modifier is for a sampled situation, the event is closed and will not reopen until the monitored condition turns false and turns true again. However, using a persistent count of 2 or more causes the event to close. But if the condition remains true for the number of intervals equal to the persistent count, the event will reopen for the same condition for which it originally fired. (See step 4 on page 271 in the procedure for defining the condition to read about persistence and how to set it.)

About this task

This procedure assumes that you are creating or editing a situation in the Situation editor (see "Opening the Situation editor" on page 255. Take these steps to add an Until modifier to the situation.

Procedure

1. Click the ⊙ Until tab.

   If it is not open already, open the Situation editor and select the situation to which you want to specify the Until modifier.

2. Specify when you want the situation to close:

   - Check **Another Situation is TRUE** to close the event when the situation is true and select a situation from the ▾ list.

     The situations available are for the same type of monitoring agent that this situation uses or that use the IBM Tivoli Monitoring common attributes.

   - Check **Reset Interval Expires** to close the event after the specified interval, 15 minutes by default.

     To change the interval, click inside dd (day), hh (hour), mm (minute) or ss (second) and click the arrows to adjust up or down, or type the number.

   See "Until" on page 266 for more information on choosing an option.

3. Select other tabs in the Situation editor to:

   - Distribute the situation to the managed systems on which you want it to run. See "Distributing the situation" on page 273.

   - Add expert advice to the situation. See "Writing expert advice" on page 275.

   - Specify an action to take or a message to send to the Universal Message Console when the situation is true. See "Specifying an action" on page 277

   - Forward events that open for this situation to an Event Integration Facility receiver such as the Netcool/OMNIbus Probe for Tivoli EIF. See "Forwarding the event to an EIF receiver" on page 280.

4. When you are finished composing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor.

   The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

   If you opened the Situation editor from the toolbar, the situation you just created will not be associated with any Navigator items and no event indicator

is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

Situations are stored at the hub monitoring server. In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error will occur, ensuring that one situation does not overwrite another.

5. If the Distribution Change window is displayed, click **OK** to add the managed systems shown in the **Required Managed Systems** list to the distribution for this situation.

Forwarding the event to an EIF receiver

When situation event forwarding has been configured in your monitoring environment, a default Event Integration Facility receiver is assigned and the severity of forwarded events matches the state specified for the situation. The EIF tab of the Situation editor enables you to specify the severity and the EIF receiver or receivers the event shall be forwarded to, or to disable event forwarding for a situation.

About this task

Note: The EIF tab shows in the Situation editor only if the hub monitoring server has been configured to forward events. See *IBM Tivoli Monitoring Installation and Setup Guide* for details.

You can avoid duplicate situation events (one from the monitoring server and one from the event server or ObjectServer) by configuring connectors as described in "Best practices for using event synchronization" in the *IBM Tivoli Monitoring Administrator's Guide*.

Complete these steps to customize how events are forwarded to the EIF receiver:

Procedure

1. In the Situation editor, click the **EIF** tab.
2. Select ☑ **Forward Events to an EIF Receiver** to send an EIF event for every event that opens for this situation or clear the check box if you prefer not to forward events that open for this situation.
3. Select the **EIF Severity** of forwarded events. <Default EIF Severity> causes the event severity that is set for the situation at this Navigator item to be applied to the EIF event when ☑ **Forward Events to an EIF Receiver** is selected.

   If the situation event being forwarded has the ℹ️ Informational severity, it will be seen as ✅ Harmless at the Tivoli Enterprise Console.
4. Specify the EIF destinations for forwarded events:
   - To add a destination, select it from the **Available EIF Receivers** list and then click the ◀ left-arrow.
   - To remove a destination from the list, select it in the **Assigned EIF Receivers** list and click the ➡ right-arrow.

   EIF destinations are defined in the Manage Tivoli Enterprise Monitoring Services application, described in *IBM Tivoli Monitoring Installation and Setup Guide*; or with the `tacmd createEventDest` command.
5. If you want to customize how situation events are mapped to EIF events, click **EIF Slot Customization**.
6. Select other tabs in the Situation editor to:

- Distribute the situation to the managed systems on which you want it to run. See "Distributing the situation" on page 273.
- Add expert advice to the situation. See "Writing expert advice" on page 275.
- Specify an action to take or a message to send to the Universal Message Console when the situation is true. See "Specifying an action" on page 277
- Specify how long a situation shall remain true. See "Adding an Until modifier" on page 279.

7. When you are finished composing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor. The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

   If you opened the Situation editor from the toolbar, the situation you just created will not be associated with any Navigator items and no event indicator is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

   Situations are stored at the hub monitoring server. In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error will occur, ensuring that one situation does not overwrite another.

8. If the Distribution Change window is displayed, click **OK** to add the managed systems shown in the **Required Managed Systems** list to the distribution for this situation.

EIF Slot Customization

Use the EIF Slot Customization editor to customize how situation events are mapped to forwarded EIF events.

The EIF Slot Customization editor opens when you click **EIF Slot Customization** in the EIF tab of the Situation editor. Here you can create a custom event mapping that overrides the default mapping between situation events and events that are forwarded to an EIF receiver such as IBM Tivoli Netcool/OMNIbus or the Tivoli Enterprise Console event server. Specify a mapping between each EIF event slot and a corresponding Tivoli Monitoring attribute, and any options that determine how attribute data is to be converted in the forwarded event.

Typically, if a mapping is not provided for a base slot, the default mapping is used, except in the case for the msg slot. However, the event mapping specified using EIF slot customization completely replaces the default mapping for extended slots. Therefore, if you do not specify a mapping for a particular extended slot or the base msg slot in this window, that slot is not included in the forwarded event. You can use the **Map all attributes** option to apply the default mapping to the base msg slot and any extended slot that is not specifically listed.

Event class name

Specify the EIF event class you want to use for the events that are forwarded to the EIF receiver. You can either type the class name directly in the field, or click **Select event class** to select from a list of all defined classes. If you do not see the name of the class you want to customize, contact your IBM Tivoli Monitoring administrator. The administrator might need to run the Tivoli Enterprise Console Event Definition Generator tool to update the list of classes. For more information on the tool, see the *IBM Tivoli Monitoring Administrator's Guide* if you are

integrating with Tivoli Enterprise Console, or see the *IBM Tivoli Monitoring Installation and Setup Guide* if you are integrating with Tivoli Netcool/OMNIbus.

When you select an event class from the list or type the name of a defined class, the **Extended slots** table is populated with the extended slots that are defined for the selected class. If you type the name of a class that is not defined in the Tivoli Monitoring configuration, the extended slots are not available and must be entered manually.

Base slots

Use the **Base slots** table to specify how each base slot of the forwarded EIF event is populated with data:
- You can specify a Tivoli Monitoring attribute to be mapped to each slot, along with options that affect how the value of the attribute is mapped to a corresponding value in the EIF event slot.
- You can specify a literal value that is used for the event slot for each forwarded event.

The list of event slots in the **Base slots** table is fixed because these slots are the same for all event classes.

Extended slots

Use the **Extended slots** table to specify how each extended slot of the forwarded EIF event is populated with data. As with the **Base slots** table, you can specify either a literal value or a mapping from an attribute.

The list of event slots in the **Extended slots** table varies depending on the extended slots defined for the EIF event class selected in the **Event class name** field. The Extended slots are included when a situation event is opened. However, they are not included in status update events when the situation event is acknowledged, unacknowledged, or closed.

If the specified class is not defined, you can manually specify the extended slots:
- To add a slot, type the slot name in any empty row of the **Slot name** column. (The new slot name must be unique; if you type a name that duplicates an existing slot name, the table cell is highlighted in red.)
- To insert an empty row in the table, right-click in any row and then click **Insert row above** or **Insert row below** from the pop-up menu. (The order of slots in the table is not significant.)
- To delete a slot, right-click in the row you want to delete and then click **Delete row** from the pop-up menu.

You can also delete an extended slot if the event definition has changed and the slot is no longer valid.

Map all attributes

Select the ☐ **Map all attributes** check box to specify that all event attributes be mapped to EIF event slots:
- If this option is selected, the default mapping is used for any base or extended slots for which you have not specified a mapping in the Base Slots table or Extended Slots table.

- If this option is not selected, all of the base slots except the msg slot are mapped by using the mapping defined in the Base Slots table or, if the base slot is not mapped, the default mapping is used. If you want to include the msg slot in events, you must map it. Also, only the slots with mappings in the Extended Slots table are included when an event is opened.

This option is useful when you want to specify a customized mapping for only a few extended slots and use the default mapping for all others.

A class name must be selected for this check box to be enabled.

Editing the mapping tables

The columns in the **Base slots** and **Extended slots** tables are as follows:

Column	Description
Slot name	The EIF event slot name.
Mapped attribute	The attribute to use for the EIF event slot. Click inside a cell and select an attribute from the group associated with the situation. **Tip:** This column is not available if you specify a value in the **Literal value** column.
Process enumerations	Whether enumerations should be resolved in the mapped slot value. Select the check box if you want to populate the slot with the resolved display value of the attribute; deselect the check box to use the actual value of the attribute. This column is available only if an attribute is selected in the **Mapped attribute** column. Click ▦ **Map enumerations** to modify an enumeration or create a new enumeration. This button is available only if the adjacent ☑ **Process enumerations** check box is selected.
Multiplier	The multiplier to apply to the attribute value when mapping it to the EIF event slot. For example, if you want to convert minutes to seconds in the EIF event, you would specify a multiplier of 60. This column is available only if a numeric attribute is selected in the **Mapped attribute** column.
Time format	The display format to use to represent a time value in the EIF event slot. Click inside a cell and select a format from the list. This column is available only if a time attribute is selected in the **Mapped attribute** column.

Column	Description
Literal value	The literal value to use for the EIF event slot. Type the value to use for this slot in all events forwarded for this situation. When the Base **Slot name** is msg, the literal value is the message template. The message template consists of fix message text and variable substitution references, or *symbols*. The symbol can refer to common or event slot data or a special reference to the situation formula. Common slots are those that are included in all forwarded events, such as situation_name; event slots are those specific to the situation. Syntax: • For an event slot, use the fully qualified attribute name ($Attribute_Table.Attribute_Name$) • For a common slot, use the variable name that is not fully qualified (without any '.') unless it is the situation symbol • For a situation formula, use *$formula$* These characters are not supported: < less than, > greater than, " quote, ' single quote, and & ampersand. This column is available only if *no* value is selected in the **Mapped attribute** column. **Note:** Variable substitution is available only for the msg slot literal value.

Note:

- For the msg slot, typical users specify a **Literal value**, not a **Mapped attribute** value. If a value is specified in the **Mapped attribute** column for the msg slot, the following results occur:
  - If **Map all attributes** is not selected, the **Mapped attribute** for the msg will not be present in the event, and will be ignored.
  - If **Map all attributes** is selected, the event includes only the default message template and not the **Mapped attribute** specified in the msg slot.
- If you select a **Mapped Attribute** for a base slot, the mapping is used when a situation event is opened. However, the mapping does not occur in situation status update events that are sent to the EIF receiver when an event is acknowledged, unacknowledged, or closed.

Map enumerations:

Use the **Map enumerations** window to modify or create enumerations used to map a Tivoli Monitoring attribute value to an EIF event slot.

The Map enumerations window is opened from the EIF Slot Customization window after you click ▦ (which does not show in that window until you select

☑ **Process enumerations** for an event slot).

If an enumeration already exists for the attribute, it populates the table and you can modify the entries. These enumerated values are used to convert attribute values to display values for the slots of the forwarded event. For example, if the

situation event uses the numerals 0 through 6 to represent days of the week, you can map those values to their names, Sunday through Saturday.

To modify the enumeration, click in the table cells and edit the text:

- Type or edit the attribute **Value** and the corresponding **Display as** value, one row for each enumeration. Any new values you add must be unique: If you type a value that duplicates an existing entry in the **Value** column, the table cell is highlighted in red.
- To insert an empty row in the table, right-click in any row and then click **Insert row above** or **Insert row below** from the pop-up menu. The order of values in the table is not significant.
- To delete a slot, right-click in the row and click **Delete row**.

Starting or stopping a situation

You can start or stop one or more situations as you find necessary for your environment.

About this task

You can start or stop one or more situations from several places: situation event flyover list; situation event console view; situation editor; or Manage Situations at Managed System. No matter which method you choose, the effect is the same: the situation starts or stops on all the managed systems to which it is distributed.

Note:
- When you stop a situation that embeds other situations, only the embedding situation is stopped; the situations that are embedded in it continue to run. This is also true for correlated situations.
- If you stop a situation that has an open event, you can close the event from the situation flyover list, the event Navigator item, or situation event console view.
- When a situation has been stopped that also has an open event that has been forwarded to the Tivoli Enterprise Console, the effect is that all open Tivoli Enterprise Console events whose cms_hostname and situation_name (as specified in the event synchronization rules) is the same as that of the stopped situation will be closed. This is done to maintain event status synchronization between the monitoring server and the event server.
- For situations within policies, you might notice an inconsistency between the situation's current status and its status displayed in the Manage Situation at Managed System window. For example, create situation A and policy B (when situation A is true, write a message to the Universal Message console). Next, stop situation A. The message log indicates a situation status of Stopped, but the Manage Situation at Managed System window indicates that the situation status is Open. In the absence of policies, the Manage Situation at Managed System window reflects the current state of a situation. However, when policies are running, this assumption is not valid. A similar inconsistency occurs if you use the situation in the UNTIL clause of another situation.

To use this feature, your user ID must have Modify permission for Situations.

Procedure

1. Depending on your location, do one of the following:

Origination point	Action
Situation event flyover list	Point to an event indicator in the Navigator to open the event flyover list, then right-click the situation name.
Manage Situations at Managed System	In the Navigator Physical view, right-click the ⊞ enterprise item or an ⊙ agent item and click ⬛ **Manage Situations**. Then, in the Manage Situations at Managed System window, select the situation that you want to start or stop.
Situation event console view	Open or create a workspace with the ⊞ **Situation Event Console** view, then right-click the row of the situation you want to start or stop.
Situation editor	Do one of the following to open the Situation editor: • Right-click a Navigator item and click ⬛ **Situations**. (See Note if you do not see this option.) • Click ⬛ **Situations** in the toolbar. • Right-click a situation in the event flyover list and click ⬛ **Edit Situation**. • Right-click an event item in the Navigator and click ⬛ **Edit Situation**. • Right-click an open event in the situation event console view and click ⬛ **Edit Situation**. • Right-click the ⊞ enterprise item or an ⊙ agent item in the Navigator Physical view and click ⬛ **Manage Situations**. Select a situation and click ⬛ **Edit Situation**. Then, right-click the name in the Situation tree on the left. (If you do not see the situation, click ⬛ **Situation Filter** and select all the options.)

2. If you want to select multiple situations to stop, use Ctrl+click to select additional situations or Shift+click to select all situations between and including the first selection and this one.

3. Click one of the following:

   ▷ **Start the situation**
   > This has no effect if the situation is already started. If a progress bar is displayed, you can click **Stop** to pause, then **Resume** when you are ready to continue, or **Cancel** to stop the operation.

   ⬤ **Stop the situation running**
   > The situation will remain stopped until you start it again manually or, if it is set to **Run at startup**, when you next start the Tivoli Enterprise Monitoring Server to which an agent connects or edit the situation.

Situation association with Navigator items

A situation must be associated with a Navigator item before you can see alerts when an event has opened.

When you create a situation, these elements of the situation definition and Navigator item properties must be set before event indicators for true situations can show on the Navigator:

Situation distribution
> The situation must be distributed to the managed systems on which it is to run. This is done in the Distribution tab of the Situation editor.

Assignment
> Situations that were created in the Situation editor after it was opened from a Navigator item, are automatically associated with the Navigator item. This is possible because the Navigator item and the managed systems assigned to it are known. Situations that were created in the Situation editor after it was opened from the main toolbar are not associated with a Navigator item initially: the situation might be distributed to one or more managed systems, but the Navigator item is not known at this point.
>
> For custom Navigator views, you must relate, or *assign*, managed systems to the Navigator items where you want to see alerts.

Association
> You must associate a situation with the Navigator items where you want to see event indicators for the situation. Association with a Navigator item is done automatically when you open the Situation editor from the Navigator, then create or edit a situation and click **Apply** or **OK**. Association with a Navigator item is done manually for situations created in the Situation editor when it was opened from the toolbar.

Tip: When you associate a situation with a Navigator item, the situation state and sound are also associated with the item. The same situation can be associated with a different Navigator item if you want to apply a different state or sound. This is especially useful for custom Navigator views where one situation might be ⊗ critical in an operator's Navigator view, but considered 🛈 informational in a view designed for an executive.

Associating a situation

If you access the Situation editor through the toolbar to create situations, no alert will be visible in the Navigator when the situation becomes true until you associate the situations with a Navigator item.

About this task

For each Navigator item you want to associate a situation with, do the following:

Procedure

1. Right-click the Navigator item to which you want to associate a situation. If this is a correlated situation, right-click the ▦ Enterprise Navigator item.
2. Click ▨ **Situations**. ▨ **Situations** is displayed in the pop-up menu for all Navigator items in the Physical view except for those at the operating platform level (such as ▨ Linux Systems). The Enterprise item is limited to Tivoli Enterprise Monitoring Server situations distributed to *HUB_ or *ALL_CMS

managed systems. A Navigator item in a custom Navigator view must have at least one managed system assigned to it, or you will not see ⬛ **Situations** in the pop-up menu.

The Situation editor opens with the list of situations that have been associated with this Navigator item.

3. Click ⬗ **Situation Filter** to open the Show Situations window.

4. Select ☑ **Eligible for Association** and clear ☐ **Associated with this Object**. If you are associating a situation with the Enterprise item in the Navigator Physical view, also select ☑ **Associated with Monitored Application**.

5. If you want to apply a different event state to the situation than the default critical state (or the one it was created with) or you want to apply an event sound, select the situation from the tree, then do one or both of the following.

   • In the **State** area, select a different state.
   • In the "Sound" area, select ☑ **Enable**. If you want to change the sound, click **Edit**.

   This ability to change the event state and sound enables you to qualify the severity of a situation for different Navigator items.

6. Right-click the situation to associate with the Navigator item and click ⇄ **Associate**.

7. When you are finished, check your work by clicking ⬗ **Situation Filter**, clearing ☐ **Eligible for Association**, and selecting ☑ **Associated with this Object**. The situations you associated with this Navigator item are displayed in the list.

8. To save your changes, click **Apply** to keep the Situation editor open; -OR- click **OK** to close it.

Results

Tip: If you have multiple situations to associate, you can use Ctrl+click to select additional situations to associate or Shift+click to select all situations between and including the first selection and this one. Be aware that the ⊗ Critical state and no sound will be applied to this Navigator item.

Dissociating a situation

If you no longer want to have event indicators originate for this situation from the Navigator item, dissociate it. For example, you might have a situation that is only embedded in other situations and never runs on its own.

About this task

To dissociate a situation:

Procedure

1. Right-click the Navigator item from which you want to dissociate a situation.

2. Click ⬛ **Situations**. The Situations option is available from all Navigator items except those at the operating platform level (such as ⬛ Linux Systems). The Enterprise item is limited to Tivoli Enterprise Monitoring Server situations distributed to *HUB_ or *ALL_CMS managed systems. A Navigator item in a custom Navigator view must have at least one managed system assigned to it; otherwise, you will not see ⬛ **Situations** in the pop-up menu.

   The Situation editor opens with the list of situations that have been associated with this Navigator item.

3. Right-click the situation you want to dissociate from the Navigator item and click 	Dissociate. If you want to remove all associations, click 	Dissociate All. If you have multiple situations to dissociate, you can use Ctrl+click to select more situations or Shift+click to select all situations between and including the first selection and this one. Then right-click one of the situations and select 	Dissociate.

4. To save your changes, click **Apply** if you want to keep the Situation editor open; -OR- click **OK** if you want to close it.

Modeling conditions for situations

Use situation modeling to perform threshold analysis on selected data samples or a situation and use the results to create or edit a situation.

You can watch the Model Situation video to learn how you can determine situation thresholds based on historical data and statistical functions.

Understanding situation modeling

Situation modeling can be done with selected attributes from the current query-based view or with the conditions of the current situation.

Modeling from a query-based view

Query-based views and historical data are valuable for exploring your monitored environment and researching conditions to build situations against. The situation modeling feature enables you to capture data from a query-based view and use it to model possible threshold scenarios or to test the effectiveness of a situation. The *Model Situation* window has options for calculating different values, comparing attributes, and transferring the results to the Situation editor to create a new situation or to fine tune the conditions in an existing situation.

If the attribute group data has been collected historically, then you can run statistical algorithms against it and provide values on the same chart. The model situation chart shows the difference between the original and calculated values. For example, you might be viewing data samples over a month-long period and see that a metric peaks during Monday mornings from 8:00 to 10:00. You can use the modeling tools to help you make decisions about whether that maximum value is repeatable or an anomaly, and whether it correlates with other conditions.

You can open the Model Situation window from the selected data row or point in a chart view or table view.

Modeling from the Manage Situations at Managed Systems window

The Manage Situations at Managed System window displays all the situations that are associated with the Navigator item for the agent. You can use situation modeling to refine the thresholds of an existing situation so that events are opened appropriately: not too often, but not so seldom that important issues are missed.

Open the Model Situation window for a selected situation to see the formula with the current attribute values plotted in the model situation chart. Use the expression editor and statistical functions to model other possible thresholds. Then copy the results into the Situation editor for editing the situation values or to create another situation.
The tool is disabled if a situation is ineligible for modeling. Situation definitions that have any these characteristics cannot be used for modeling:

- Embedded situations, which also includes correlated situations
- Situations with more than ten attributes in the formula
- Situations with no numeric attributes in the formula
- Situations with an attribute repeated in the formula (such as MyAttribute > 0 AND MyAttribute < 50)
- Situations with Boolean OR logic
- Situations with group functions: count, sum, average, minimum, maximum

Modeling a situation with view data

Use the values from the current chart or table to model threshold conditions and then use the results in a new situation. Historical views are particularly valuable for researching thresholds.

Before you begin

Your user ID must have View permissions for Situations for this feature to be enabled. To create or edit Situations, your user ID must have Modify permission.

About this task

The attribute values that are used for modeling are especially valuable for comparing with historical data samples. Take these steps to work with the situation modeling feature and, if you choose to, create a situation. Text attributes cannot be modeled. Only display item, which has a checked box

Procedure

1. Right-click a table row or chart point containing numeric attributes to monitor and click ![icon] **Model Situation**. The Select attribute window is opened with the attribute group preselected. If you opened the pop-up menu from a chart, attributes are preselected.

2. In the **Select attribute** window, select the attributes to model:
   - Click an attribute to select it. You can also use Ctrl+click to select multiple items and Shift+click to select all items between the item last selected and this one. You can select up to 10 attributes.
   - Select numeric attributes of interest.

   After you click **OK**, the Model Situation window is opened with a chart displaying the selected numeric attributes and their values (from the selected row).
   - If historical data samples are available, the last 24 hours are displayed in a plot chart. If no historical data has been collected, the current data samples are displayed in a bar chart.
   - Any non-numeric attribute values are shown but are non-selectable. The exception is a display item, which is a key attribute in some multiple row attribute groups and shows in the Model Situation window. In situation modeling, it can be used to select different data rows. In a situation definition, a display item can be selected as an **Advanced** option that enables the situation to open multiple events: once for each row that meets the condition.
   - A message informs you if no modeling chart is displayed. This happens if the situation override is for a multiple-row attribute group that has no display item.

3. Use the modeling tools to apply different threshold scenarios; click **Recalculate** to render the chart with the modified argument set.

4. To create a situation populated with values from the Condition column, click **Create Situation**. Enter the situation name and optional description.

Related reference:

"Model Situation" on page 292

You can select ⊞ **Model Situation** from a query-based view or Manage Situations at Managed System window to display selected attribute values in a chart, try out different thresholding scenarios, and use the results to create or edit a situation.

Modeling with an existing situation

Use situation modeling to calculate new threshold values for an existing situation or to create another situation with the modeling results, replacing the original situation thresholds. This enables you to fine tune the situation so that it does not open events unnecessarily or, conversely, miss important issues.

Before you begin

To get the most benefit from situation modeling, use historical data samples for comparison. If no historical data is being collected for the attribute in the situation, you can still model thresholds but without the benefit of historical data and statistical functions.

Some situation constructs make them ineligible for modeling: embedded situations, and that includes correlated situations; more than ten attributes in the formula; no numeric attributes in the formula; attributes repeated in the formula; and Boolean OR logic.

Your user ID must have Modify permissions for Situations for this feature to be enabled.

About this task

Take these steps to model threshold conditions for a situation, then apply the results to the situation or to a new situation.

Procedure

1. Right-click a Navigator item that has situations associated with it: either ⊞ Enterprise or ◉ Agent.

2. Click ⊞ **Manage Situations**.

3. Click the row of the situation that you want to model and click ⊞ **Model Situation**. The Model situation window opens with the formula in the **Attributes** expression editor and the current values plotted in a chart. If historical data samples are available, the last 24 hours are plotted. Any non-numeric attribute values are shown but are non-selectable.

4. Use the modeling tools to try out different threshold scenarios; click **Recalculate** to render the chart with the modified argument set.

5. To create a situation populated with values from the Condition column, click **Create Situation**; or to edit a situation, click **Edit Situation**.

Related reference:

"Model Situation"
You can select ⊞ **Model Situation** from a query-based view or Manage Situations at Managed System window to display selected attribute values in a chart, try out different thresholding scenarios, and use the results to create or edit a situation.

Model Situation

You can select ⊞ **Model Situation** from a query-based view or Manage Situations at Managed System window to display selected attribute values in a chart, try out different thresholding scenarios, and use the results to create or edit a situation.

If the attribute group used by the query-based view or situation has historical data collection enabled, the Model Situation window shows a plot chart in the top half of the window (up to the last 24 hours) and Attributes, Time Span, and Functions areas in the lower half. You can switch the display to a bar chart or area chart. If historical collection has not been configured, the Model Situation window shows a bar chart in the top half of the window and an Attributes area in the bottom half.

Model situation chart

If you started from a query-based view, your selections from the Select Attribute window are displayed in a chart with values from the current row or data point. If you started from a situation, the formula is displayed in the Attributes expression editor and the current attribute values are plotted in a chart. As you modify attribute conditions (such as to change the value), the chart is adjusted and the difference indicated.

For attributes with historical data samples, the last 24 hours of data samples are plotted by default unless you have a time span set, in which case that time span is plotted. Although all the data points are plotted, the window has only limited space for time stamp labels. Use any of these methods to see values and adjust the display:

- Move the mouse pointer over a point to see the value and time stamp.
- Drag the window border to expand the chart width.
- Drag the border bar between the chart and the Attributes area up or down.
- Click 🔍 **Zoom in** (Ctrl + Click) and drag the mouse over the time period to highlight; the selection expands to fill the chart area. Click 🔍 **Zoom out** (Esc) to return to the full data view.
- Select options in the **Time Span** area to reduce the time span and thus the number of samples displayed.
- Click 📊 , 🗺 , or 🖼 to redraw the data as a bar chart, plot chart (the default), or area chart.

Attributes

This area has a row for each selected attribute or situation attribute. Click inside a cell to see only that attribute plotted on the chart.

Name   This is the name of the attribute chosen from the Select attribute dialog.

If you see a ☑ check box in the first column, it means the attribute is a *display item*, a key attribute in a multiple-row attribute group. You can enter another valid value in the Condition cell to plot the values for that display item attribute.

Condition

This is the value of the attribute from the chart view data series or table view row, or from the situation if you are modeling with an existing situation.

Click inside a `▼ > 85816   ▼` cell to change the function or operator, or to edit the value. A solid marker line is drawn across the chart to represent the current comparison value. A shaded region appears above or below this marker line if the relational operator represents a range of trigger values (>, >=, <, <=). The difference between the original and calculated value is the *threshold area* and is highlighted.

Value modifier

This is initially set to < none selected >. Click inside the cell to calculate a threshold by raising or lowering the original value by a specified percentage.

In the + PERCENT or - PERCENT window that opens, enter the percentage to raise or lower the value by, click **Draw** to apply the threshold to the chart. After you click **OK**, the difference between the original and percentage raised or lowered is used to set a new threshold value.

✐ Chart options

Modify the appearance of the threshold: ☐ **Visible in legend** to add **Threshold** to the legend, ☐ **Show chart label** to add a threshold label, and **Fill Style** to change the threshold pattern or color. Click **Custom Colors** to choose from a wider selection of colors or to create your own.

Functions

This area is displayed when the attribute group used in the view or situation has historically collected values. Click a row in the **Attributes** area. In the **Functions** area, select the check box for the statistical function you want to use. After setting an **Argument**, click **Recalculate** to see the adjusted threshold in the **Result** cell and plot it on the chart as a colored line.

☐ **AVG**

For the selected attribute, calculate the average value from the samples taken during the time period, and add or subtract the standard deviation that was entered in the **Argument** field. Valid values are -6 to 6.

☐ **MIN**

For the selected attribute, find the lowest sampled value within the time period, and increase or decrease that value by the percentage that was entered in the **Argument** field. Valid values are -100 to 100.

☐ **MAX**

For the selected attribute, find the highest sampled value within the time period and increase or decrease that value by the percentage that was entered in the **Argument** field. Valid values are -100 to 100.

☐ **PERCENTILE**

For the selected attribute, find the value below which NUM percent of the data points fall; where NUM is the value from the **Argument** field. Valid values are 1 to 100.

☐ **MODE**

For the selected attribute, find the most frequently observed value within the time period. No argument is entered for a MODE calculation.

⌨ **Copy**

> Add the value from the **Result** cell to the selected cell in the Attributes area.

✎ **Chart options**

> Modify the appearance of the threshold: ☐ **Visible in legend** to add **Threshold** to the legend, ☐ **Show chart label** to add a threshold label, and **Fill Style** to change the threshold pattern or color. Click **Custom Colors** to choose from a wider selection of colors or to create your own.

Time Span

This area is displayed when the attribute group used in the view or situation has historically collected values. The **Last** 24 Hours is shown by default.

Last   You can change from the default 24 hours to a different value or unit of time (Hours, Days, Weeks, Months, Quarters, Years).

> The **Time Column** field shows the attribute for the time stamp, such as **Recording Time**.

Custom

> Use the **Custom** tab to plot a specific date range or to select a different time attribute.

> Click inside the **Start Time** and **End Time** fields to select a date and time.

> The **Time Column** field shows the attribute for the time stamp, such as **Recording Time**.

Daily   Use the **Daily** tab to plot a specific time range.

> Click inside the **Start Time** and **End Time** fields to select a date and time.

> Select the ☑ **Show data for last** ☐ check box and specify the number of days to plot.

> Select ◉ **calendar day** to include any weekend days that fall within the date range; or ◯ **weekdays** to exclude weekend days. For example, if today is Monday and you select the last 2 weekdays, Thursday and Friday of the past week are plotted.

> Select the ☐ **Show this day's data** check box to select a specific date to plot within the time range.

Buttons

Recalculate

> Click **Recalculate** after you have entered an argument for a statistical function to see the result in the **Result** column and plotted on the chart. Use this button also to refresh the chart with any changes you make to the time span selection or override key conditions.

OK   Click **OK** to use the values shown in the Attributes area.

Create Situation

> This button opens the Create Situation window for you to create a new situation. The situation is created immediately after you type a name and click **OK**. The Situation editor opens with the modeling attributes and conditions in the formula editor. Continue to refine the formula and complete the situation definition.

Reset   After you click inside an attribute cell, the chart is redrawn to plot only

that attribute. If you also select a statistical function and click **Recalculate**, the statistical result is also plotted. Click **Reset** to redraw the chart with all attributes.

These are all the reset actions taken:
- Cancel any attribute selection in the table
- Clear any shaded threshold area in the chart
- Show all chart series if any were not visible
- Remove all statistical markers from the chart
- Clear the Function fields, including function results
- Cancel any zooms in progress
- Disable the create schedule action

Cancel

Click Cancel to exit the dialog without saving your changes.

Example

This example of situation modeling was started from the **Total % Privileged Time** and **% User Time** stacked bar in the Process Time Distribution chart from the Windows OS Processor workspace.

Initial chart display

We right-click a bar in the chart and select **Model Situation**. The current data sample is 8% privileged time and 4% user time. Looking at the historical samples, we see that values over 70% and 40% are rare.

Change % Privileged Time threshold

We adjust the % Privileged Time threshold by clicking in the **Condition** column and changing from **> 8** to **> 70**, which is reflected in the chart as the gray patterned area.

Change % User Time threshold

Next, we adjust the % User Time threshold by clicking in the **Condition** column and changing from **> 4** to **> 40**, which is reflected in the chart as the gray patterned area.

The thresholds give a visual indicator of how often a situation created with these values would become true.

Create Situation

After we click **Create Situation**, the thresholds are inserted as expressions in the Situation editor. With the situation formula as it is now, an event will be opened when privileged time is over 70% and, at the same time, user time is over 40%.

	% Privileged Time	% User Time
1	> 70	> 40

Edit the situation formula

We moved the % User Time value to a new row so that the situation becomes true when *either* the privileged time is over 70% or the user time is over 40%. Be aware that now that the situation has Boolean OR logic, it is not eligible for situation modeling; thus, when the situation is selected in the Manage Situations at Managed System window, ⊞ **Model Situation** will be disabled.

	% Privileged Time	% User Time
1	> 70	
2		> 40

Related tasks:

"Modeling a situation with view data" on page 290
Use the values from the current chart or table to model threshold conditions and then use the results in a new situation. Historical views are particularly valuable for researching thresholds.

"Modeling with an existing situation" on page 291
Use situation modeling to calculate new threshold values for an existing situation or to create another situation with the modeling results, replacing the original situation thresholds. This enables you to fine tune the situation so that it does not open events unnecessarily or, conversely, miss important issues.

Editing a situation

Situation maintenance is done in the Situation editor. You might want to fine tune a situation so it fires when the condition is truly at a warning or critical level and not too early or too late. Or you might want to delete situations you no longer need.

About this task

Note: If you plan to modify a predefined situation, copy the situation with ⬤ **Create Another** and edit the copy. This ensures that you can start again with the original should you need to.

👤 Your user ID must have Modify Situation permissions to open the Situation editor for creating and maintaining situations.

Follow these steps to edit a situation:

1. Open the Situation editor using one of the access points described in "Opening the Situation editor" on page 255.

   The Situation tree lists the situations for the level chosen: at the system level, you will see all situations applied to that system; at the agent level, you will see all situations written with attributes from that agent; and at the attribute level, you will see all situations written with attributes from the associated group or groups.

2. Click a situation to open it.

   If you do not see the situation in the tree and you opened the editor from the Navigator menu, click 🢒 **Situation Filter** and select ☑ **Eligible for Association**.

3. Click a tab to modify that part of the definition:

   - 𝑓𝑥 Formula tab to edit the condition definition, sampling interval, or status.

     Tip: If you need to add or remove rows or clear cell contents, right-click and select one of the pop-up menu options.
   - 🖻 Distribution tab to see and assign the systems to monitor
   - ⓘ Expert Advice tab to enter text viewable from the event results workspace
   - ⚡ Action tab to enter a command to be invoked when the situation fires
   - 🕓 Until tab to have the event close after a period of time or when another situation is true

4. When you are finished editing the situation, click **Apply** to save your changes and start the situation; -OR- **OK** to save your changes, start the situation, and exit the Situation editor.

   The Situation editor performs syntax checking on each expression you enter. If you enter an improper value for an attribute, such as "30000" for % User Time, the cell turns red and the situation will not be saved until you correct the error.

   When the formula or other aspects of a situation definition are changed, the situation is first stopped and then re-started. An event indicator for the situation will appear in the Navigator for a brief period and the event flyover will show the ⬤ situation stop indicator.

Deleting a situation

Use the Situation editor to delete situations. Keep in mind that deleting a situation removes it from every managed system to which it was distributed.

About this task

To delete a situation:

Procedure

Do one of the following to open the Situation editor:

- Right-click a Navigator item and click ⊞ **Situations**. (See Note if you do not see this option.)
- Click ⊞ in the toolbar.
- Right-click a situation in the event flyover list and click 📝 **Edit Situation**.
- Right-click an event item in the Navigator and click 📝 **Edit Situation**.
- Right-click an open event in the situation event console view and click 📝 **Edit Situation**.
- Right-click the ⊞ enterprise item or an ⊚ agent item in the Navigator Physical view and click ⊞ **Manage Situations**. Select a situation and click 📝 **Edit Situation**.

Results

1. Right-click the situation name in the tree on the left.

   If you do not see the situation, click ⚷ **Situation Filter** and select all of the options.

   You can use Ctrl+click to select additional situations or Shift+click to select all situations between and including the first selection and this one.

2. Click ✖ **Delete** to permanently remove the situation selection.

   If a progress bar is displayed, you can click **Stop** to pause, then **Resume** when you are ready to continue, or **Cancel** to stop the operation.

3. If this situation is referenced (embedded) in another situation or in a policy, a Referential Integrity window opens with a list of the affected objects and you must do the following before you can delete the situation:

   a. Click 📋 **Copy** if you want to save the list to the Windows clipboard.

   b. Click **Cancel**.

      You can now paste the list into a text editor or a notepad view in a workspace.

   c. Edit each situation or policy in the list to delete the reference to this situation. (With the situation open in the Formula editor, right-click the column heading for the embedded situation, and click ✖ **Delete**. With the policy open in the Workflow editor, select the activity references this situation, and click ✖ **Clear**.)

Note: ⊞ **Situations** is only in the pop-up menu of Navigator items that have managed systems assigned. In the Navigator Physical view, managed systems are preassigned to every item except those at the operating platform level (such as Linux Systems) and cannot be changed. In the Navigator Logical and custom views, you can assign managed systems when creating a Navigator item or editing its properties.

Note: If you are forwarding events to an event integration facility such as the Tivoli Enterprise Console Server, when an active situation is deleted, the Stop and Delete events that are sent to the event server use the situation short name and not the long name. This is also true for any associated messages that are sent to the message log view.

Situation overrides for dynamic thresholding

Use situation overrides for dynamic thresholding at selected managed systems, under specific conditions, and at scheduled times.

Overriding a situation formula is useful if you need to adjust the comparison values based on local conditions or a specific time. The values used in a situation override can be either fixed or calculated from runtime data. Examples include the following:

- Overriding the formula for a particular managed system because of its unique configuration. For example, you might want to use a different threshold to detect a "disk full" condition for a system that has a large hard disk drive.
- Focusing on the performance of a specific process or component. For example, you might override a comparison value only when the Process Name attribute matches a specified value, in order to specify a lower CPU time threshold for that process.
- Applying different expression values during specified periods of time to accommodate varying service level objectives. For example, you might apply different override expressions during weekends and holidays.
- Deriving comparison values from historical data for a managed system. For example, you might specify that the threshold for a "high CPU usage" situation should be based on deviation from the average CPU usage for that system.

Creating a situation override

Use the situation expression override function to change the formula thresholds for one or more of the managed systems that the situation is distributed to. The override can remain in effect indefinitely, at a set schedule, or until you remove it.

Before you begin

When you define an override, apply it immediately or at a scheduled time. The schedule for an override can run on specific days (*calendar schedule*) or every day at a certain time (*hourly schedule*). Calendar schedules can be saved and selected at any time. An hourly schedule can be assigned to the schedule, but is deleted after you close the Override editor if you turn off the assignment. If you plan to schedule the override, decide whether you want a calendar schedule or hourly.

Make sure the situation is eligible for overriding, and that it is currently distributed to one or more managed systems either directly or as part of managed system groups. If a situation is not eligible, **Override Formula** remains disabled even after you select a managed system or managed system group.

About this task

The Expression Override editor is available through the **Distribution** tab of the Situation editor.

Procedure

1. Identify the managed system or managed systems where the override should take place and click **Override Formula** from either of these windows:
   - On the Situation editor **Distribution** tab, select one or more managed systems or managed system groups from the **Assigned** list. The Override Formula button is disabled until you select an entry from the Assigned list.
   - In the Manage Situations at Managed System window, select the situation that you want to override and click *fx* **Override Formula**.

   The Add Expression Overrides window opens.
2. If the Expressions tab has no expressions, click **Add Expression** to add the reference formula.

3. Edit the expression value (or more if there are multiple expressions in the formula) for the overrides to apply. These values will be used instead of those in the reference formula when the override definition is applied.

- Click inside a cell and enter a new value. After you press Enter or click outside the cell, the new value appears with a blue background to indicate that the original reference value has changed.

- Click **Model Expression** if you want to use expression modeling to establish threshold values based on historical data and statistical analysis. See "Model Expression" on page 309.

- The use of wildcards (* or ?) cannot be used in the key value.

Examples of eligible repeated attribute expressions (situations involving these patterns are eligible for override because the operators are unique):
`%_Processor_Time >= 30 and %_Processor_Time <= 50 p_name == "Websphere"`

4. To view or change the override details, click the **Details** tab. The details include the description, distribution, and override priority (1 is the highest priority; 200 is the lowest).

5. If you want to apply the override immediately and without defining a schedule (unless you defined one in the Model Expression editor), click **OK** and skip to step 7. Otherwise, continue to the next step.

6. Specify a schedule for the override:
   a. Select ▦ **Click here to assign an override schedule**.
   b. In the Select Override Expression Schedule dialog, either select one of the defined schedules entries and continue to step 7 or click ▯ **Create New Schedule**. Any schedules created from the command line with **tacmd addCalendarEntry** are displayed in the list.
   c. Select ◉ **Create a new calendar schedule** to create a schedule for particular days or sequence of days, or ○ **Create a new hourly schedule** to create a daily schedule that runs at specified times.
   d. In the Schedule editor that opens, define the schedule: "Calendar" on page 306 or "Hourly" on page 308

   After you click **OK**, the undefined schedule icon changes to indicate that a ▦ calendar schedule or an ▦ hourly schedule has been applied.

7. Click **Apply** to validate the override definition and save your changes. The situation override is validated when you save it. If any problems are found, a message describes the error. Select the ☑ check box next to the ▦ Calendar error and click **Remove Expression**.

8. If you want to add another override, click **Add Expression** and repeat these steps starting at step 3 A copy of the reference formula is added to the formula overrides. You can have one formula that has no schedule; any others must have a schedule applied.

Results

If the situation override entry has no schedule, it begins running after you click **Apply** or **OK**. Any overrides with a schedule are activated at the designated time.

What to do next

You can use the monitored baselines feature to see the override expressions charted. See Adding monitored baselines to a chart.

Outside of the portal, you can use the CLI **tacmd setOverride**, **tacmd addCalendarEntry**, **tacmd suggestBaseline**, and **tacmd acceptBaseline** commands to create override expressions based on calculated baseline values rather than fixed values. (described in the *IBM Tivoli Monitoring Command Reference*) or using a local thresholds.xml configuration file (described in the "Enterprise situation override XML specification" in the *IBM Tivoli Monitoring Administrator's Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/adminuse/itm_admin.htm)).

If the **Calendars** list shows only **<No schedule>** or only the calendar entries that you have defined and no predefined calendars, it indicates that the hub monitoring server was upgraded from an earlier version but application support was not added. Application support is added by running the CLI `itmcmd support` command. See the *IBM Tivoli Monitoring Installation and Setup Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/install/itm_install.htm) for instructions.

Adaptive monitoring considerations

Review these considerations when you are planning to implement adaptive monitoring, such as which situations are eligible for expression overrides.

To create a situation override, you must be editing a situation that does not have any of these characteristics

- Situations that embed other situations.
- Correlated situations. These are run at the hub monitoring server.
- Situations that are distributed only to the monitoring server.
- Situations containing multiple attribute groups.
- Situations containing expressions that involve aggregate column functions: MIN, MAX, AVG, SUM, COUNT, CHANGE, PCTCHANG.
- Situations containing expressions involving these row functions: STR, SCAN, MISSING, TIME, DATE.
- Situations containing multiple expressions using the same attribute *and* the same relational operator (for example, `p_state != "running"` and `p_state != "waiting"`).

Override Formula is not available in the Situation editor's **Distribution** tab or in the Manage Situations at Managed System window if the situation is not eligible for expression overrides.

The following restrictions apply to override expressions

- An override expression cannot add attribute comparisons that do not appear in the reference formula or change the comparison operators used.
- If the override contains multiple expressions, each must have a unique override condition.
- An override expression must be different from the reference formula (at least one comparison value must be changed).
- If you no longer need an expression override definition, you can remove it (in the Add Expression Overrides window, select the check box next to the definition and click **Remove Expression**). After clicking **OK** to close the Add Expression Overrides window, you will be unable to edit the situation until it has been refreshed: You can adjust the value in an expression, but you cannot change the formula function or relational operator.

The following restrictions apply to override expression modeling
- The modeling graph is unavailable if the situation override is for a managed system group.
- The modeling graph is unavailable if the situation override is for a multiple-row attribute group that has no display item (used as a key condition).

Avoid conflicting or overlapping override expressions

For a situation that has multiple expression overrides defined for the same managed system or managed system group, avoid defining conflicting or overlapping override expressions. A typical conflict scenario is where two or more override expressions have schedules assigned and an overlap exists in their time ranges. In these cases the order in which the individual override expressions are evaluated by the target agent or agents is undefined.

Keep the situation override expressions under 4000 bytes

The total number of characters used in all the expression overrides defined for a situation should not exceed 4000 bytes. The actual size requirement for a single override varies depending on the names and values of the key columns and the override expression. In one case the limit might be 25 or, in a simpler case, it might be higher.

The symptom of exceeding the 4000-byte limit is that the overrides do not work and the monitoring server trace log shows an "exceeds limit 4000" override error. See the ITM Situation Limits > Situation Formula Overrides technote for details.

Situation overrides defined at the command line

You can also use the Command Line Interface **tacmd setOverride** to define a situation override and **tacmd addCalendarEntry** to create a calendar schedule for the override.

A more granular schedule can be defined with **tacmd addCalendarEntry**, which uses the cron format, than with the Tivoli Enterprise Portal Schedule editor. For example, in the Schedule editor you can specify a calendar schedule for the override that runs once every hour, once every 2 hours, and so on, in 1-hour increments (such as every 3 hours at 00:00, 03:00, 06:00, and so on). In the CLI you can specify interval times to the minute, such as `-c 10 * * * *` for every 10 minutes. If you create a schedule from the CLI that is too complex for the Schedule editor, an error message is displayed and you must use the CLI to modify the schedule.

See the *IBM Tivoli Monitoring Command Reference* for details about the tacmds.

Disabling and enabling situation overrides

A situation override starts as soon as you define it or when you have it scheduled. You can also manually enable and disable situation overrides. This procedure applies to any override definition for a situation, scheduled or not.

Before you begin

This procedure enables or disables all the overrides for the situation. If you have any overrides that you no longer need to keep, you can selectively remove them beforehand.

Procedure

1. Click ▣ **Situation Editor** or right-click the Navigator item that the situation is associated with and click ▣ **Situations**.

2. In the Situations tree, right-click the situation. Situations with override definitions show an *ƒx* override indicator at the right end of the status bar.

3. Click one of the following options:

   - ⦿ **Enable Situation Overrides** to activate the override (or overrides) for the situation. Any unscheduled and scheduled overrides for this situation are now enabled. Unscheduled overrides begin running immediately; scheduled overrides are activated at their designated time.

   - ⊖ **Disable Situation Overrides** to deactivate the override (or overrides) for the situation. Any expression overrides for this situation that are running are stopped and disabled; the original situation definition is now in effect. Any scheduled overrides do not resume until you have enabled them again manually or edit the schedule.

Removing override definitions

When you no longer want to keep a situation override, remove the definition. You can remove a specific override definition or all the override definitions that have been written for a situation.

Before you begin

This procedure permanently deletes the selected override definition (or definitions). If you think you might want to use the override definition again at a later time, you can disable it rather than remove it.

About this task

Take one of these steps to remove the selected override definition or to remove all the override definition for a situation.

Procedure

- Remove an override definition for a situation:
  1. Open the Add Expression Overrides editor:
     - In the Situation editor, open the situation that has an override definition assigned (indicated by *ƒx* at the right end of the status bar). In the Distribution tab, select the managed system (or managed systems) from the **Assigned** list and click **Override Formula**.
     - In the Manage Situations at Managed System window, select the situation that has an override definition assigned (indicated by *ƒx* in the Overrides Exist column) and click *ƒx* **Override Formula**.
  2. Select the check box for each override definition that you want to remove and click **Remove**.
  3. Save your changes by clicking **Apply** to keep the window open or **OK** to close the window.

- Remove all the override definitions for a situation:
  1. Click ▣ **Situation Editor** or right-click the Navigator item that the situation is associated with and click ▣ **Situations**.
  2. In the Situations tree, right-click the situation. (Use Ctrl + click to select multiple entries or Shift + click to select all entries between the one

previously selected and this one.) Situations with override definitions show an *fx* override indicator at the right end of the status bar.

3. Click ⊠ **Remove All Overrides**.

What to do next

If you are in the Situation editor and want to edit the definition, you first must refresh the situation. To do this, click another situation in the tree, then click this situation again. After the situation has been refreshed, you no longer see *fx* in the status bar.

Expression override and scheduling windows

Use the Expression override editor to create overrides for dynamic thresholding and use the override schedule dialogs to determine when they should run.

Add Expression Overrides

You can use the Add Expression Overrides window to define a situation override to be distributed to one or more managed systems.

A situation override is a set of expressions that can conditionally substitute comparison values that are different from those in the situation formula, and applied to one or more of the managed systems that the situation is distributed to.

The Add Expression Overrides window supports creating expression overrides using specified comparison values. You can also use expression modeling to create the override expression from a calculated baseline value.

Navigation tree

The navigation tree on the left side of the window shows the situation you are currently editing and the assigned managed systems and managed system groups for the situation, with the currently selected entries highlighted. You can click items in this list to change which systems are included in the override.

Reference formula

The **Reference formula** area shows the original situation formula. This is the formula you are overriding. You can view the reference formula as a table or graphic by clicking the respective ▦ **Table View** or 🔆 **Graphic View** tab.

Formula overrides

The **Expressions** tab is divided into two panes: on the left are attributes that can be applied to the override, but which cannot be used for expression modeling; on the right are attributes whose override values can be modified with expression modeling. Specify the override values to apply to the selected managed systems.

- To model override values, click **Model Expression** (see "Model Expression" on page 309). This feature is most effective for attributes that have historical data collected)

- To edit an expression, click in a table cell and edit the values.

  – To define a schedule when the override should run, select ▦ **Click here to assign an override schedule**. Select a defined schedule from the list or create a new schedule (see "Select Override Expression Schedule" on page 305).

  – The columns in the right-hand section define the comparison values for the expression. These are the values that will be substituted for

those in the reference formula if the override condition is met, or if no override condition is set. Click in the appropriate table cells to edit the comparison values used for the override expression.

- To add an expression, click **Add Expression**.
- To remove an expression, select the check box next to the expression that you want to remove and click **Remove Expression**.

The **Details** tab shows information for the situation override:

Assigned to
> The managed systems or managed system groups to which this situation override is currently distributed.

Priority

> This field is a display-only field showing the relative priority of this situation override. If multiple overrides are distributed to the same managed system, the priority determines which override takes precedence. This can happen if an override is distributed to an individual managed system, and another override is distributed to a managed system group that includes the same system.

> Situation overrides that are distributed to managed systems have a higher priority than overrides distributed to managed systems groups: 1 is the highest priority; 200 is the lowest.

Description
> Enter an optional description of the situation override.

Add Expression
> Click this button to add a new formula row.

Remove Expression
> Select the check box next to the expression you want to delete and click **Remove Expression**.

Model Expression
> Click this button to open the Model Expression dialog.

Restriction:
- An override expression cannot add attribute comparisons that do not appear in the reference formula or change the comparison operators used.
- If the override contains multiple expressions, each must have a unique override condition.
- An override expression must be different from the reference formula (at least one comparison value must be changed).
- If you no longer need an expression override definition, you can remove it (in the Add Expression Overrides window, select the check box next to the definition and click **Remove Expression**). After clicking **OK** to close the Add Expression Overrides window, you will be unable to edit the situation until it has been refreshed: You can adjust the value in an expression, but you cannot change the formula function or relational operator.

Select Override Expression Schedule

Use the **Select Override Expression Schedule** dialog to select a calendar entry that specifies when the expression override should be applied. You can define a new schedule in this dialog and edit or delete existing schedules.

From the **Add Expression Overrides** dialog, click ▦ **Calendar** to establish a schedule, ▦ **Calendar** to edit a calendar schedule, or ▦ **Calendar** to edit an hourly schedule.

Schedules is a list of schedules that are available for this situation. Click a schedule to see its description in the **Description** field. The predefined schedules are NonPrimeShift, PrimeShift, Weekday, and Weekend. Use the following tools to create, edit, or delete a schedule:

Create New Schedule
> Define a new schedule including name, description, and timing.

Create Another Schedule
> Copy the selected schedule and give it a new name. This is a good choice if you want to base the new schedule on, or in contrast to, what was defined for the source schedule.

Modify Schedule
> Edit a schedule to change the time or frequency that it runs.

Delete Schedule
> Delete a schedule that you no longer need for this situation.

If you want to keep a schedule definition but disable it, click **Remove**. The schedule is deselected and **<No schedule>** is enabled.

Schedule Options

Use the Schedule Options to select whether to create a new calendar schedule or hourly schedule. The new schedule options are displayed from the Select Override Expression Schedule window after you click ▯ **Create New Schedule**.

Create a new calendar schedule
> Choose this option to create a schedule for a particular day and time or to repeat at a set frequency. A calendar schedule is useful if you want the override to run only once and starting at a particular time or if you want it to repeat at set intervals.

Create a new hourly schedule
> Choose this option to create a daily schedule that runs at specified times. An hourly schedule is desirable when you want the override to run every day, starting and stopping at a particular time from 00:00 to 23:45.

Schedule editor

Use the Schedule editor to create or modify a schedule to control when the situation override is applied. The editor is opened from the Select Override Expression Schedule dialog after you click ▯ **Create New Schedule** and select the schedule type (Calendar or Hourly) or after you click ▨ **Modify Schedule**.

The Schedule editor has some different options for and schedules.

Calendar

You can have the override run only once or repeat on the specified days.

Schedule identity
> The **Name** and **Description** fields are used to identify the schedule. The name entered here is shown in the list of schedules in the Select Override Expression Schedule dialog. The description is shown only here and can be used to enter information such as the purpose of the schedule.

Schedule definition

Select from the ⏷ list whether to apply the schedule **once** or at intervals. Your choice changes the options in the rest of the definition area:

Frequency

once   Select the month from the **On month** ⏷ list and any value from 1 to 28, 30, or 31 from the **On day** ⏷ list.

daily   Select a value from 1 to 30 from the **Every (number of days)** ⏷ list.

weekly

Select the check box for the day to apply the schedule. On days of the week: ☐ All days or a specific day or selection of days: ☐ Sun ☐ Mon ☐ Tue ☐ Wed ☐ Thu ☐ Fri ☐ Sat

monthly by day

Select the frequency of months from the **Every (number of months)** ⏷ list. Select the corresponding check box for the day or days to apply the schedule.

monthly by date

Select the frequency of months from the **Every (number of months)** ⏷ list. Select the corresponding check box for the day or days of the month to apply the schedule.

◯ **Run at time** ☐

Runs the override once at the selected time for every selected day. Select a time in 15-minute increments from 00:00 to 23:45 or click inside the field and type the time (such as 23:10).

◯ **Run all day**

Runs the override every minute of every hour starting at 00:00 of the day specified, with the last sampling at 23:59. This option is disabled if the schedule is set to run **once**.

◯ **Run every** [1 ⬍] **hours**

Runs once on the hour and repeats at the selected hourly interval. For example, Run every 4 hours runs the override at 00:00, 04:00, 08:00, 12:00, 16:00, and 20:00. This option is disabled if the schedule is set to run **once**.

◯ **Run during these ranges**

Runs every minute from the starting hour to one minute before the stopping hour. This option is disabled if the schedule is set to run **once**.

Select the radio button to enable the **From** and **Through** fields. Then click inside a field to select the start and stop times from 00:00 to 23:59.

Click **Add Range** to add another range row for specifying multiple time ranges. Select a check box and click **Remove Range** to delete the range.

The following example shows a schedule that applies the expression override from 06:00 to 07:59 and from 20:00 to 21:59.

Because the check box for range 2 is selected, we can click **Remove Range** to delete range 2.

Note: The command line interface **tacmd addCalendarEntry** enables you to create more granular schedules than in the Tivoli Enterprise Portal using the cron format. If you create a schedule from the CLI with a syntax that is not editable in the Schedule editor, such as `-c 5 20-21 * * *` to run every 5 minutes from 20:00 to 22:00, an error message is displayed and you must use the CLI to modify the schedule.

Schedule preview
This is a calendar of the days that the situation override takes place with a yellow highlight and the days that are skipped with a blue highlight. In the following example, the days of the month that have already passed have no highlight. The override is scheduled to occur every three days.

Hourly

With an hourly schedule, the override takes place every day starting and stopping at the time indicated. Hourly schedules are temporary: After creating an hourly schedule, if you do not assign it to the override, it is removed from the Schedules list after you close the Select Override Expression Schedule dialog box.

Schedule identity
The **Name** and **Description** fields are disabled because they are filled in for you from the start and stop times that you enter in the definition. For example, after selecting 16:00 and 21:00 for the start and end times, you get *16:00-21:00* as the name and *Daily schedule between 16:00 and 21:00* as the description.

Schedule definition
Click the **Start time** field and select a time (in quarter-hour increments) from the list or type the start time yourself following the **hh:mm** format shown. Repeat for the **Stop time** field.

Schedule preview
This is a calendar of the days that the situation override takes place highlighted with yellow. The days of the month that have already passed have no highlight.

	2010					>>
<	January					>

SUN	MON	TUE	WED	THU	FRI	SAT
*	*	*	*	*	1	2
3	4	5	6	7	8	9
10	11	12	13	14	15	16
17	18	19	20	21	22	23
24	25	26	27	28	29	30
31	*	*	*	*	*	*

After you click **OK** to save the schedule, the schedule name is displayed in the
Schedules list of the Select Override Expression Schedule dialog.

Model Expression

Select the Model Expression button in the Add Expression Overrides dialog to use
visual baselining for testing threshold scenarios and use the results to create an
expression override for the situation.

Use any of these methods to see values and adjust the display:
- Move the mouse pointer over a point to see the value and time stamp.
- Drag the window border to expand the chart width.
- Drag the border bar between the chart and the Attributes area up or down.
- Click ⊕ **Zoom in** (Ctrl + Click) and drag the mouse over the time period to
  highlight; the selection expands to fill the chart area. Click ⊖ **Zoom out** (Esc) to
  return to the full data view.
- Select options in the **Time Span** area to reduce the time span and thus the
  number of samples displayed.
- Click ▊, ▧, or ▨ to redraw the data as a bar chart, plot chart (the default), or
  area chart.

Expressions

This area shows the values of the current expression override, one row for each
expression. Initially, these are the same values used in the reference formula (the
original situation formula).

☐ ▦   If you want to delete an expression, select the check box and click **Remove
Expression**. If you delete all the expressions, click **Add Expression** to add
the reference formula.

    You can specify a schedule condition for an expression in two ways:
- Select ▦ **Click here to assign an override schedule**. Select a defined
  schedule from the list or create a new schedule (see "Select Override
  Expression Schedule" on page 305).
- Click ▦ the **Select an area in the chart to create a schedule for** in the
  toolbar. Drag the mouse from left to right or from right to left to
  highlight the time range when the schedule should run. After you
  release the mouse button, a new expression row is added with a daily
  schedule defined for that time period, which is expressed on the chart as
  a shaded area. (See example.)

    A defined schedule is indicated by ▦ for a schedule that runs on specific
days or ▦ for a schedule that runs daily within the specified time period.

Column headings
    These are the names of the attributes referenced in the situation. Any
attributes on the left side of the divider are designated as display items for

the situation and can be used as an override condition. The override expression is applied only when the comparisons you specify are true. The attribute on the right side of the divider are the comparison values for the expression being modeled. These values are used instead of those in the reference formula when the override condition is true, or if no override condition is set.

Click inside a comparison value cell on the right side of the divider to see only that attribute plotted on the chart and to edit the value.

After you edit a value and press Enter, the cell has a ■ blue background to indicate it has changed from the original value. A line is drawn across the chart to represent the current comparison value. A shaded region appears above or below this marker line if the relational operator represents a range of trigger values (>, >=, <, <=).

If you click inside a blue cell and change the value, it is plotted in the chart automatically. This enables you to fine tune a threshold value and observe the affected values in the chart.

If you click another object (such as a Time Span or Function field) while an attribute cell is active, the cell is highlighted in ■ yellow. If you click a point in the chart area, the value where you clicked is inserted in the cell.

Functions

This area is displayed when the attribute group used in the situation has historically collected values. Click a cell in the **Expressions** area to enable the **Functions** area, and select the check box for the statistical function to use. Optionally enter a value in the **Argument** field and click **Recalculate** to see the adjusted threshold in the **Result** field and plot it on the chart as a colored line.

☐ **AVG**

    For the selected attribute, calculate the average value from the samples taken during the time period, and add or subtract the standard deviation that was entered in the **Argument** field. Valid values are -6 to 6.

☐ **MIN**

    For the selected attribute, find the lowest sampled value within the time period, and increase or decrease that value by the percentage that was entered in the **Argument** field. Valid values are -100 to 100.

☐ **MAX**

    For the selected attribute, find the highest sampled value within the time period and increase or decrease that value by the percentage that was entered in the **Argument** field. Valid values are -100 to 100.

☐ **PERCENTILE**

    For the selected attribute, find the value below which NUM percent of the data points fall; where NUM is the value from the **Argument** field. Valid values are 1 to 100.

☐ **MODE**

    For the selected attribute, find the most frequently observed value within the time period. No argument is entered for a MODE calculation.

Copy

    Add the value from the **Result** cell to the selected cell in the Attributes area.

☑ **Chart options**
Modify the appearance of the threshold: ☐ **Visible in legend** to add **Threshold** to the legend, ☐ **Show chart label** to add a threshold label, and **Fill Style** to change the threshold pattern or color. Click **Custom Colors** to choose from a wider selection of colors or to create your own.

Time Span

This area is displayed when the attribute group used in the situation has historically collected values. The **Last** 24 Hours is shown by default.

Last   You can change from the default 24 hours to a different value or unit of time (Hours, Days, Weeks, Months, Quarters, Years).

The **Time Column** field shows the attribute for the time stamp, such as **Recording Time**.

Custom
Use the **Custom** tab to plot a specific date range or to select a different time attribute.

Click inside the **Start Time** and **End Time** fields to select a date and time.

The **Time Column** field shows the attribute for the time stamp, such as **Recording Time**.

Daily   Use the **Daily** tab to plot a specific time range.

Click inside the **Start Time** and **End Time** fields to select a date and time.

Select the ☑ **Show data for last** ☐ check box and specify the number of days to plot.

Select ◉ **calendar day** to include any weekend days that fall within the date range; or ○ **weekdays** to exclude weekend days. For example, if today is Monday and you select the last 2 weekdays, Thursday and Friday of the past week are plotted.

Select the ☐ **Show this day's data** check box to select a specific date to plot within the time range.

Buttons

Recalculate
Click **Recalculate** after you have entered an argument for a statistical function to see the result in the **Result** column and plotted on the chart. Use this button also to refresh the chart with any changes you make to the time span selection or override key conditions.

OK   Click **OK** to close the dialog and use the values from the Expressions area for the expression override.

Reset   After you click inside an attribute cell, the chart is redrawn to plot only that attribute. If you also select a statistical function and click **Recalculate**, the statistical result is also plotted. Click **Reset** to redraw the chart with all attributes.

These are all the reset actions taken:
- Cancel any attribute selection in the table
- Clear any shaded threshold area in the chart
- Show all chart series if any were not visible
- Remove all statistical markers from the chart

- Clear the Function fields, including function results
- Cancel any zooms in progress
- Disable the create schedule action

Cancel

Click **Cancel** to exit the dialog without saving your changes.

Example

This example shows a daily schedule over three days that applies the override >= 75 from 01:00 to 06:00.

Embedding a situation

Reference a situation in another situation by embedding it. Embedding situations enables you to monitor several conditions with the same situation and mix and match them to create new situations, rather than having to manually enter different combinations of expressions.

Before you begin

To use this feature, your user ID must have Modify permission for Situations

About this task

You embed a situation by selecting ⦿ **Situation Comparison** for **Condition Type** in the Select Condition window. The Select Condition window appears after you click **OK** on the Create Situation window, or after you select **Add Condition...** on the **Formula** tab. Take these steps to embed a situation:

Procedure

1. In the Select Condition, for the Condition Type, select ⦿ **Situation Comparison**.
2. Select the situation to embed. You can then use Ctrl+click to select additional situations from the list or Shift+click to select all situations between and including the first selection and this one. Situations that have already been chosen are colored blue.
3. Click **OK** to see the situations in the formula editor.
4. Click in a cell below each new column to add ▦ **== True** and complete the expression. If two conditions must be satisfied, this is an AND comparison and their expressions must be entered on the same row of the editor; if either

condition can be satisfied for the situation to be true, then this is an OR comparison and the expressions are entered on different rows.

5. A situation that embeds another requires at least two conditions. If you have more conditions to add, click **Add Conditions**, and repeat steps 1 through 4 until you are finished with the formula. The other condition can be a formula function or another embedded situation.

6. When you are finished composing or editing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor.

Results

The Situation editor performs syntax checking on each expression you enter. If you entered an improper value for an attribute, such as "abc" for Disk Size, the cell turns red and the situation will not be saved until you correct the error.

The properties of the embedding situation, such as the monitoring interval, override the properties of the embedded situation.

In the unlikely event that two users attempt to save a situation with the same name simultaneously, an error will occur ensuring that one situation does not overwrite another.

What to do next

If you opened the Situation editor from the toolbar, the situation you just created will not be associated with any Navigator items and no event indicator is displayed when the situation becomes true. You must associate a situation with a Navigator item or change its state and click **Apply**.

You must distribute each embedded situation to the same managed systems as the embedding situation. Otherwise, the embedding situation does not run.

If you change an embedded situation, you must restart each situation that embeds it.

When you stop a situation that embeds other situations, only the embedding situation is stopped; the embedded situations continue to run.

A situation cannot be embedded in the situation if it has a display item assigned. You can create another situation and remove the display item assignment from the new situation to enable it for embedding.

Situations are stored at the hub monitoring server. It is possible that while you were building the situation that any of the embedded situations could have be deleted. Should this happen, a message will tell you that the embedded situation is missing when you attempt to save your changes. You will need to remove the embedded situation from this one before you can save it. You can always create a new situation to be embedded or find another one that is the equivalent and embed that one.

Correlating situations

Create a correlated situation, which is made up entirely of other situations, to target a set of conditions on a subset of the original distribution. For example, you might have two situations that signify a critical breakdown when they both become true on a particular managed system. You can create a correlated situation of these two situations and distribute it to the managed system in question.

Before you begin

To use this feature, your user ID must have Modify permission for Situations and have the Enterprise Monitoring Server as an Allowed Application.

About this task

The behavior of situations embedded in a correlated situation can each be targeted to monitor a subset of their original distribution.

Suppose you have a Disk_Usage_Warning situation that has been distributed to 100 managed systems. The correlated situation in this examples alerts you when Disk Space Warning is true on the most critical managed systems:

	Disk_Usage_Warning @Primary:MAIN1:NT	Disk_Usage_Warning @Primary:MAIN2:NT	Disk_Usage_Warning @Primary:MAIN3:NT
1	== True	== True	== True

This correlated situation can be edited so that it becomes true when the Disk_Usage_Warning situation opens an event on any two of these three managed systems:

	Disk_Usage_Warning @Primary:MAIN1:NT	Disk_Usage_Warning @Primary:MAIN2:NT	Disk_Usage_Warning @Primary:MAIN3:NT
1	== True	== True	
2	== True		== True
3		== True	== True

Take these steps to create a correlated situation:

Procedure

1. Do one of the following to open the Situation editor:
   - Click ⬚ if, for the time being, you do not want to see an event indicator when this situation becomes true. (You can associate the situation at a later time.)
   - Right-click the ⬚ Enterprise Navigator item and click ⬚ **Situations**.

   Correlated situations are run at the Hub Tivoli Enterprise Monitoring Server, so they can be associated only with the Enterprise Navigator item. If this is a custom Navigator view, the association can be to any Navigator item where the hub (HUB_HOSTNAME) is assigned.
2. Click ⬚ **New Situation** to open the Create Situation window.
3. Type a name for the situation in the **Name** field and, optionally, a description in the **Description** field. The situation name can contain any UTF-8 characters

and can be up to 256 bytes long. The description can be up to 64 characters long. The description is displayed in the Situation editor, in the Manage Situations window, and when you select the situation for embedding in or correlating with another situation.

4. Select `Correlated Situation` from the **Type** list and click **OK**. After you select Correlated Situation, the **Monitored Application** field is set to `All Managed Systems` and disabled. Correlated situations are distributed to the hub monitoring server automatically.

5. In the Correlated Situation Selection window, select a situation from the **Available Situations** list, select the managed system (or managed systems) to correlate from the **Available Managed Systems** list, then click **OK** to add the situation to the editor. A column is inserted for each managed system that was selected.

6. Click in a cell below each new column to add ☑ **== True** and complete the expression. If the comparison must be satisfied on both managed systems, this is an AND comparison and their expressions must be entered on the same row; if the comparison can be satisfied on either managed system, this is an OR comparison and the expressions must be entered on different rows.

7. A correlated situation is not complete until it has at least two conditions. If you have more situations to add, click **Add Conditions**, select the next situation and managed system (or managed systems), click **OK**, and select inside a cell of each column added. If a managed system name is colored blue instead of black, it means the situation for this managed system was already added to the correlation formula.

8. Adjust any other settings in the **Formula** tab to complete the formula:
   - To add event persistence, click **Advanced**.
   - To change the sampling interval, click inside the dd (day), hh (hour), mm (minute) or ss (second) field and click the up and down arrows.
   - To start the situation manually rather than automatically, clear ☑ .
   - To choose a different **State**, select one from the list. (Not available if you opened the Situation editor from the main toolbar.)
   - To play a sound when an event opens, select ☑ in the **Sound** area. Click **Edit** to change the sound file or its properties. (Not available if you opened the Situation editor from the main toolbar.)

9. Specify options on the other tabs as needed.

10. When you are finished composing or editing the situation, save your changes and start the situation by clicking **Apply** to keep the editor open; or by clicking **OK** to close the Situation editor.

Results

The properties of the embedding situation, such as the monitoring interval, override the properties of the embedded situation.

If you enter a system command in the Action tab, the **Attribute Substitution** button is disabled because a correlated situation references one or more situations and not attribute groups.

When you stop a situation that embeds other situations, only the embedding situation is stopped; the embedded situations continue to run.

What to do next

If you change an embedded situation, you must restart each situation that embeds it.

If you opened the Situation editor from the toolbar or from any Navigator item other than the one to which the hub monitoring server is assigned, the situation you just created will not be associated with any Navigator items and no event indicator is displayed when the situation becomes true. If this is the case, right-click the ⊞ Enterprise Navigator item and click ▨ **Situations**. In the Situation editor, click 🐾 **Situation Filter** and select ☑ **Eligible for Association**. In the Situation tree, right-click the correlated situation and click ⇄ **Associate**, then**Apply** or **OK**.

A situation cannot be referenced in the correlated situation if it has a display item assigned. You can create another situation and remove the display item assignment from the new situation to enable it for inclusion in a correlated situation.

Tutorial: Correlating situations across managed resources

You can gets hands on experience in correlating situations across managed resources by completing the tasks in this tutorial.

Tasks to accomplish the goal

The following tasks will accomplish the goal:

1. Create a situation that monitors for the missing process and opens an Informational situation event.
2. Create a correlated situation to reference this situation and to become true for any two of the three computers. When the situation becomes true it will open a Warning situation event.
3. Create a second correlated situation much like the other, but have it become true only when the process is missing on all three computers. When the situation becomes true it will open a Critical situation event.

👤 To use this feature, your user ID must have Modify permission for Situations.

Note: The following scenario assumes that you have the Windows OS monitoring agent installed.

The situations developed for your monitoring environment can be designed to convey the true urgency of the situation event while ensuring that operators are not overwhelmed by unnecessary alerts.

In this scenario, a company has software used by the support staff that must be running on any of three computers at all times. The goal is to find out when the software is not running, and to prioritize the alerts so that when the process is stopped on two systems, a warning alert is issued; and if it is stopped on all three systems, a critical alert is issued.

First task

Create a situation that monitors for the missing process and opens an Informational situation event:

1. Right-click the Process (**Enterprise > Windows Systems > Windows OS > Process**) Navigator item and click ▨ **Situations**.

2. In the Situation editor, click ⬜ **New Situation** to open the Create Situation window.

3. Type `Missing Process` in the **Name** field (where `Process` is the name of the executable software), select `Windows OS` as the **Monitored Application**, and click **OK**.

4. In the Select Condition window, select the NT_Process attribute group and the Process Name attribute, then click **OK**.

5. Click inside the row 1 Process Name cell, click the 🔽 **Function** field and select 🔍 Check for Missing Items.

6. Type the name of the software process in the Missing Item List and click **OK**.

   For the hands-on practice, type notepad.

7. Select ℹ️ **Informational** from the State field.

8. Click the **Distribution** tab, move the other two managed systems to the **Assigned** list.

   The Assigned list must show only the names of the three managed systems to be monitored for this software.

9. Click **Apply**.

The situation is now running on all the managed systems of interest and will provide an informational alert if the software stops on any of the managed systems.

Tip: If you were not planning to use this in a correlated situation, you could have chosen Process_Name for the Display Item (**Advanced** button in the **Formula** tab), which enables a situation event to be opened for each managed system where the process is missing rather than only the first one found.

Second task

Create a correlated situation to reference this situation and to become true for any two of the three computers. When the situation becomes true it will open a Warning situation event

1. While the Situation editor is open, click ⬜ **New Situation**.

2. Type a name for the correlated situation (such as No_Notepad_Warning).

3. Select `Correlated Situation` **Type** and click **OK**.

4. In the Correlated Situation Selection window that opens, select the missing process situation (such as Missing_Notepad) from the Available Situations list.

5. Select the first two of the three managed systems from the Available Managed System Group and click **OK**.

6. Click inside each row 1 cell to add 🔲 **== True** and complete the condition.

7. Click **Add conditions**, select the first and third managed systems and click **OK**, then click inside the row 2 cells for these managed systems.

8. Click **Add conditions**, select the second and third managed systems and click **OK**, then click inside the row 3 cells for these managed systems.

   When you are finished, the formula editor should look like this:

Table 8. Formula Editor

Header	Missing_Notepad @Primary:Ali	Missing_Notepad @Primary:Beth	Missing_Notepad @Primary:Carol
1	==True	==True	
2	==True		==True
3		==True	==True

9. Select ⚠ **Warning** from the State field, and click **Apply**.

Now, as well as the original situation, you have a correlated situation that escalates the alert to Warning when the software is stopped on two managed systems.

Third task

Create a second correlated situation much like the other, but have it become true only when the process is missing on all three computers. When the situation becomes true it will open a Critical situation event

1.

With the Situation editor is open, click 📄 **Create another Situation**.

2.

Type a name for the correlated situation (such as No_Notepad_Critical).

3.

For rows 2 and 3, right click the row number and select Delete from the pop-up menu.

4. Click inside the third cell in row 1 to add the managed system to the situation definition.

The formula should look like this:

Table 9. Formula

Header	Missing_Notepad @Primary:Ali	Missing_Notepad @Primary:Beth	Missing_Notepad @Primary:Carol
1	== True	== True	== True

5. Select ⊗ **Critical** from the State field, and click **OK**.

Test the situation

Now test the scenario:

1.

Start the software (such as Notepad) on all three managed systems.

2.

Stop the software on one managed system and wait for an informational event indicator to appear on the Navigator item.

3.

Stop the software on one of the other managed systems and wait for a warning event indicator to appear.

4.

Stop the software on the third managed system and wait for the critical event indicator.

The Navigator would first show an 🔲 Informational event indicator over one of the Process Navigator items, then a 🔲 Warning indicator over two of the items and, finally, a 🔲 Critical indicator over all three items.

- 🏢 Enterprise
  - 📁 Windows Systems
  - 📁 Ali
    - ◉ Windows OS
      - 🖥 Disk
      - 🖥 Enterprise Services
      - 🖥 Memory
      - 🖥 Network
      - 🖥 Printer
      - 🔲 Process
  - 📁 Beth
    - ◉ Windows OS
      - 🖥 Disk
      - 🖥 Enterprise Services
      - 🖥 Memory
      - 🖥 Network
      - 🖥 Printer
      - 🖥 Process
  - 📁 Carol
    - ◉ Windows OS
      - 🖥 Disk
      - 🖥 Enterprise Services
      - 🖥 Memory
      - 🖥 Network
      - 🖥 Printer
      - 🖥 Process

Situation examples

Review the situation definition examples to better understand the different types of situations you can create.

Display item examples

Some multiple-row attribute groups have attributes that can be designated as display items in situations. When you add a display item to the situation, it behaves a little differently. Like any other situation, it opens an event when its condition is true. But unlike other situations, it can become true again for another row in the data sampling if the condition is true.

Processes situation

To illustrate, consider a situation that looks for any Windows process that is consuming over 5% of processor time and with Process Name designated as the display item. As it is displayed in the Formula tab of the Situation editor:

	% Processor Time
1	> %5

and in Show Formula

`(% Processor Time > 5) Display Item:NT_Process.Process_Name`

After starting the situation on MYSYSTEM, we learn that three processes are consuming more than 5%. The situation becomes true for BIG, BIGGER, and BIGGEST, and three events are opened. If this situation had no display item, it would have become true for BIG, but never would have found BIGGER or BIGGEST until the first process went below 6% and the event closed.

If you were to look at the events in the message log view, they would look like this:

Status	Situation Name	Display Item	Origin Node	Global Timestamp
▣ Opened	Process Time GT 5 percent	Biggest	Primary:SYSTEM:NT	12/15/09 11:10:05
▣ Opened	Process Time GT 5 percent	Bigger	Primary:SYSTEM:NT	12/15/09 11:09:27
▣ Opened	Process Time GT 5 percent	Big	Primary:SYSTEM:NT	12/15/09 11:09:05

Event log situation

In this example, the log space low situation checks the Windows logs for any that are over 95% capacity, and has the event log name as its display item.

	% Usage
1	>= 95

and in Show Formula:

`(% Usage >= 95) Display Item:NT_Monitored_Logs_Report.Log_Name`

After starting the situation on MYSYSTEM, we learn that two log files are almost full, Application and System, so two events are opened. If this situation had no display item, it would have become true for System, but never would have found Application until the System log file size had been reduced to less than 95% capacity.

If you were to look at the events in the message log view, they would look like this:

Status	Situation Name	Display Item	Origin Node	Global Timestamp
▣ Opened	NT Log Space Low	Application	Primary:SYSTEM:NT	12/15/09 11:10:05
▣ Opened	NT Log Space Low	System	Primary:SYSTEM:NT	12/15/09 11:09:27

Embedded situation examples

Conditions that can be combined in many different situation scenarios make useful embedded situations. The following situation can be embedded in any situation that you do not want to run on Sundays when the computer is being serviced.

Not_Sunday situation:

	Day of Week
1	>'Sunday'

NT_Services_AutoStart_NoSunday

	Start Type	Current State	Not_Sunday
1	=='Automatic'	== 'Stopped'	== True

The Not_Sunday situation is distributed to, but not associated with the managed system so you do not see an event indicator in the Navigator when it becomes true, which would otherwise happen every day except Sunday. The sampling interval for this situation was set to 1 day. This enables the situations in which it is embedded, like NT_Services_AutoStart_NoSunday, to continue running but not cause alerts on Sunday when the computer is unavailable.

Correlated situation examples

Suppose you have a Disk_Usage_Warning situation that has been distributed to 100 managed systems. You create a correlated situation that alerts you when Disk Space Warning is true on the most critical managed systems.

	Disk_Usage_Warning @Primary:MAIN1:NT	Disk_Usage_Warning @Primary:MAIN2:NT	Disk_Usage_Warning @Primary:MAIN3:NT
1	== True	== True	== True

This correlated situation can be edited so that it becomes true when the Disk_Usage_Warning situation opens an event on any two of these three managed systems:

	Disk_Usage_Warning @Primary:MAIN1:NT	Disk_Usage_Warning @Primary:MAIN2:NT	Disk_Usage_Warning @Primary:MAIN3:NT
1	== True	== True	== True
2	== True		== True
3		== True	== True

Expert advice HTML examples

```
node=$EVENT:ATTRIBUTE.ISITSTSH.NODE$;
"BODY BGCOLOR='#FFF80'><div align=center><h3><i>Advice on situation
</i></h3><h2>"+name+"</h1><p><img src=http://localhost:15200/wws_globe.gif>
The number of concurrent connections associated with the Tracker server on
<b>"+node+"</b> is high.<p> Click here for more information: <a
href=http://localhost:15200/classes/candle/fw/resources/help/
```

Situation troubleshooting

Use troubleshooting situations to resolve some problems in the environment.

As well as following these diagnostic steps, you can use a set of Web Services tools for Tivoli Monitoring that help you perform audits of the environment and correct common issues. See the IBM Integrated Service Management Library for more information.

You can also use the ITMSUPER Tools to learn about the health of your managed systems, situations, and environment configuration. The ITMSUPER Tools are included in the IBM Support Assistant (ISA), a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. To install the ISA software, go to IBM Support Assistant.

Agent not connected to the monitoring server

After a situation has been distributed to an agent and started, it continues to run until it is stopped manually. Most situations at the agent also stop when the agent to which the situations are distributed loses contact with its Tivoli Enterprise Monitoring Server.

The situations that continue to run at the agent are considered to be autonomous situations.

There are two categories of autonomous agent situations:

- Historical data collection (or UADVISOR) situations where historical data is configured to be stored at the agent. UADVISOR situations are used to gather data samples at intervals and then insert the samples in historical data files located at either the monitoring server (TEMS) or agent (TEMA). When historical data collection is stored at the agent, the UADVISOR situations continue to run when the agent loses communication with its monitoring server.

- A situation that specifies an action command (reflex automation) that is to be executed at the agent. A situation containing an action command creates two separate requests that are sent to the agent. One request returns the event data back to the monitoring server. The second request executes the command. Only the request executing the command at the agent will remain active when the monitoring server connection is lost.

Troubleshooting events not opening

A Critical, Warning, or Informational status light goes on in the Navigator when a situation associated with an item becomes true.

About this task

If a light is not displayed when you think it should be, check each item below in sequence until the issue is resolved:

Procedure

1. *Is the managed system online ?* Situations cannot run on an offline managed system.

- If the agent is offline, restart it. The installation guide for your IBM Tivoli Monitoring product has instructions for configuring and starting an agent. If you cannot restart the agent, contact IBM Support.
- If you see a status discrepancy between the Navigator and the Managed Systems Status table (one view shows the agent is offline and the other shows it online), try recycling the Tivoli Enterprise Portal Server. If the condition does not clear, contact IBM Support.

2. *Is the situation distributed to the managed system and associated with a Navigator item?* To find out, right-click the Navigator item where the event light should be displayed, then select Situations from the pop-up menu. The left frame of the Situation editor shows all situations distributed to and associated with the Navigator item. If the situation is listed, it has been properly distributed and associated.

- If you do not see the situation, click ⇲ **Situation Filter**, adjust the settings as shown below and click **OK**:

☑ **Associated with Monitored Application**

☑ **Eligible for Association**

☐ **Associated with this object**

For each IBM Tivoli Monitoring product applicable to this Navigator item, you will see:

⊙ Agent type

⧈   Distributed to this managed system but not associated with the Navigator item.

    You need to assign the Informational, Warning or Critical state and any sound, and click **Apply**.

▬   Neither distributed to this managed system nor associated with the Navigator item.

    Select the Distribution tab and distribute the situation. Then assign the Informational, Warning or Critical state and any sound , and click **Apply**.

- If you still do not see the situation, contact IBM Support.

3. *Does the message log view **show that the situation has started?*** Click the Enterprise item in the Navigator Physical view to see the message log view (or you can create one). The situation must show as Started and Open. If not, click ⧉ Situation Editor and do one of the following:

Status	Situation Name
⚠ Problem	Test_Situation
▷ Started	My_Situation
◧ Open	My_Situation

- If you do not see the situation as Started, right-click the situation name and click ▷ **Start the situation**, and then click **OK**.
- If the situation is started but no event opens, check that the condition is true. Continue to Step ▷ .
- If the message log shows the situation has a Problem, the cause might be that the formula has a syntax error and you must follow step ❹ ; -OR- there is a version mismatch between the agent support at the Tivoli Enterprise Monitoring Server and the Tivoli Enterprise Portal Server. Contact Customer

Support. This can easily happen when the Tivoli Enterprise Monitoring Server and Tivoli Enterprise Portal Server are on different computers and you upgrade the Tivoli Enterprise Portal Server and not the monitoring server.

4. *Has the condition specified in the formula been met?* It is possible that everything about the situation, its distribution and association are correct but the situation does not fire because the condition has not been met.

   Before you open the Situation editor, it is a good idea to create a table view of the same attributes as those in the situation. You can then compare the returned values against those in the formula.

   - In the Situation editor, click **Show Formula** and review the expressions to confirm that the syntax and threshold values accurately reflect the condition you are testing for.

   - Has persistence been added to the situation and has the threshold been exceeded? The light will go on and the message log will show Open after the situation has remained true through the consecutive number of intervals specified.

   - Does the situation have an UNTIL clause and is the other situation true? When the other situation is true, the event for this situation is closed even if the condition still exists. An event will not be opened until the next interval and only if the condition is true.

   - If you still can't determine the cause of the situation not firing, try adding an OR expression for one of the attributes that is sure to be true. Then click **Apply** to restart the situation with the true condition. It should become true. Remove the OR expression when you are done.

5. Contact customer support If you went through all these steps and still think it should be firing, then turn on RAS traces on the monitoring server to which the managed system is connected as instructed by Customer Support.

Troubleshooting events not closing

If you have a situation that seems always to be true and have open events, you can adjust the formula or the interval or add persistence to the definition.

About this task

Procedure

1. Are the condition thresholds too stringent? The threshold values in the formula might be appropriate to signal an alert on one managed system but not another. Try adjusting the formula to raise the threshold. Or, if the threshold should vary for the different types of managed systems you are monitoring, create another situation, edit the values, and distribute to only the managed systems to which it applies. You will also need to remove these managed systems from the distribute list of the original situation.

2. Is the interval too short? If the situation is always true, the interval might be so short as to not allow the condition to resolve itself before the next data sampling. Try increasing the interval between data samplings to a time that enables the condition to return to normal status; -OR- consider adding persistence or an UNTIL clause, described next.

3. Has persistence or an UNTIL clause been added? When the other situation is true, the event for this situation is closed even if the condition still exists. An event will not be opened until the next interval and only if the condition is true.

4. Contact customer support If you went through these steps and the situation is still true too does not work as you expect, contact Customer Support.

   Note: If you only want to stop events opening for a certain period of time while you take down a system or application for maintenance, you can disassociate situations from Navigator items. The situations will continue to run but will not send event notifications to the Navigator until you associate them again. See Associate Situations. Alternatively, you can add an UNTIL clause that closes the event when another situation is true. The other situation specifies a time span, such as Sunday after 10:00 PM or Monday before 4:00 AM. See Create a Situation; Situation Editor - Until.

High CPU usage on the monitoring server

A poorly written situation can have a surprisingly large impact on CPU usage on a monitoring server.

About this task

Using wildcard characters (*) in situations is one of the most common mistakes made in writing situations. Using fields longer than required can also cause additional work to evaluate the situation. Making the changes below have, in some instances, reduced CPU usage on a monitoring server from close to full utilization to nearly completely idle. The CPU usage manifests itself in the kdsmain process.

You might notice that kdsmain uses large amounts of CPU, agents appear offline, agent deployments fail, and other symptoms occur. Follow the guidelines for situations below.

Procedure

1. Avoid using wildcard characters in situations whenever possible.
2. Use the simplest field possible that will provide the required value. Look at the attributes under consideration in the Tivoli Enterprise Portal prior to writing the situation. There might be a field that gives the required value without performing a search operation on that field.
3. If a situation requires searching for a phrase within a field, use **Scan for string within a string**.
4. When performing a **Count of group members** operation, use a short field for the count. The field selected does not change the result, but does cause less data to be handled to derive it.
5. Avoid using Unicode fields unless they are required. They require more processing then their non-Unicode counterparts.

Example

Scenario 1: Situation to ensure syslogd is running.

   Original situation: `MISSING(Process Command (Unicode)) == ( '*syslogd*' )`

   Revised situation: `MISSING(Base Command) == ( 'syslogd' )`

The revised situation is better for several reasons. First, it removes the wildcard characters from the situation. Second, it uses a non-Unicode field. Third, it does not require any searching operation as the **Base Command** field returns exactly the required phrase.

Scenario 2: Situation checks for free space on multiple mount points starting with /opt/IBM.

Original situation: `Mount Point == '/opt/IBM*' AND Space Available Percent < 5`

Revised situation: `SCAN(Mount Point) == '/opt/IBM' AND Space Available Percent < 5`

The revised situation removed the wildcard character and used the more efficient SCAN function.

Scenario 3: Situation checks for more than 3 tecad_logfile processes.

Original Situation: `Process Command (Unicode) == *tecad_logfile* AND COUNT(Process Command (Unicode)) > '3'`

Revised Situation: `Base Command == 'tecad_logfile' AND COUNT(Base Command) > '3'`

The revised situation uses **Base Command** instead of a search or wildcard characters to find the required value. It also uses a shorter field to perform the count operation and avoids Unicode fields.

Chapter 11. Automating actions and responses

The Tivoli Enterprise Portal has functionality that enables you to design your own Take Action commands, Launch Application definitions, and Policy definitions, then store them for use by other users. A policy, rather than being run by the user, runs automatically when a referenced situation event is opened.

Automation with policies

A *policy* is a logical expression that describes a series, or *workflow*, of automated steps, called *activities*, whose order of execution you control. You can create policies that implement fully automated workflow strategies and policies that require user intervention.

Policy workflow

A policy begins by waiting for a situation to become true, then starts a series of actions. The simplest policy will simply wait until a situation is true and then take some action. For example, the predefined Windows OS policy named *NT Disk Busy* is activated when **1** the *NT Percent Disk Time* situation becomes true (disk busy more than 95%). Then **2** *NT Missing Process* situation is evaluated. If the situation is true **3** , a script is run at 1 A.M.; if the situation is false **4** , the `net start` command is run to start that service; if it is in error **5** , a script is run at 1 A.M.

Workflow activities

A policy consists of activities that can automate responses to events or routine operator tasks. An activity can be an activity program, a situation, or another policy. The workflow elements, activities and +⟶+ connectors, are represented in the Grapher view in the Workflow window much as they would be in a flow chart: Some activities flow directly to the next activity, while others branch off depending on the result of the previous activity.

Policies can have these types of activities:

- Predecessor activities, which directly precede other activities; successor activities, which directly follow other activities. They act on the basis of the *endcode* of their predecessor.
- Parallel activities that run at the same time.
- Sequential activities that run one at a time in succession.

Workflow endcodes

When an activity ends, it returns an endcode. Endcodes control the flow of activities in a policy.

Policy distribution

Policies are distributed to the managed systems that you want to monitor

and send commands to and controlled by the Tivoli Enterprise Monitoring Server that the monitoring agents connect to.

A policy starts after you save it if ☑ **Auto start** is enabled (the default setting). A policy also starts after you select it and click ▶ **Start Policy** in the Workflows editor or Manage Policies at Managed Systems window.

If the policy distribution is changed manually or dynamically, the policy is restarted. Dynamic updates happen after a new agent comes online or after an agent connection switches from one monitoring server to another.

Predefined policies

Some monitoring agents provide predefined policies that you can use as is or modify for your environment. See the user's guide for your product for a description of any predefined policies.

Viewing policy details

The Policy Details area of the Workflows window is a listing of all the policies on your managed network and their settings.

Before you begin

If the Workflows tool is disabled, your user ID does not have View Policy permission or Workspace Author Mode permission or both.

About this task

Take one of these actions to see a list of the policies that are available on your managed network or on a particular managed system.

Procedure

- Click ⚙ **Workflow Editor**. With this method, you can see and select all policies, distributed or not distributed. The Workflows window opens with the policy details at the top. If any of the column headings are obscured, you can move the mouse pointer over them to see their titles or drag the column border left or right to adjust their width.
- In the Navigator Physical view, right-click the ▦ Enterprise Navigator item or an ◉ Agent Navigator item and click ▤ Manage Policies from the popup menu. With this method, you can see and select from all the policies that are distributed to the selected managed system.

Creating a policy

Use policies to perform actions, schedule work, and automate manual tasks on one or more managed systems in your enterprise.

Before you begin

Your user ID must have Modify Policy permission. Otherwise, the Workflows tool is disabled.

About this task

The first step to building a policy for automation is to create a new policy and fill in important details such as where it is distributed to and whether it is auto started. Take these steps to begin building a new policy.

Procedure

1. Click 🧩 **Workflow Editor**.

2. Click 🗋 **New Policy** (or Alt+🗋 to keep the full Policies Details area open). New_Policy is added to the Policy Details list and a blank Grapher view opens in the Workflow editor for you to add and connect activities. To enlarge the grapher view, click ▲ Collapse, which collapses the Policy Details area and expands the Workflow Editor area to full window; click ▼ Expand to restore the Policies Details area. You can also drag the border up or down.

3. In the newly added row of the **Policy Details** area, edit the name and description, distribute the policy, and adjust the settings as needed. (See "Policy list" on page 336.)

4. In the Workflow editor, add a 🔲 **Wait until a situation is True** activity:**Wait until a situation is True** in the General activities tab of the Workflow components frame. The policy must begin with this activity type unless it is uncorrelated, in which case you can start with any activity.

   a. Click the 🔲 **Wait until a situation is True** activity.

   b. Click inside the Grapher view to place the activity icon. The Select a Situation window opens for you to select a situation that, when true or in error, triggers the next activity in the workflow.

   c. Select the situation and click **OK**. (If you do not see the situation in the list, click **Cancel**, and create the situation you need, then repeat these steps.) The activity icon displays with the name of the situation.

5. Add other activities to the policy (see Workflow Components). Double-click an activity tool to keep it on. You can add activities to start or to stop other policies or situations, to suspend execution for a given period of time, to evaluate a situation, and to invoke specific actions You can have several activity programs running either sequentially or in parallel.

6. Draw connecting lines to direct the workflow:

   a. Click ⊢═⊣ Connector, then click inside one activity and again inside the one that follows. The Select a link condition window opens with a list of possible conditions that will run the next activity. When you join activities with the Connect tool, you assign an endcode to each connection. An endcode indicates the result of activity processing and control the flow of activities in a policy.

   b. Select the result that will trigger the next activity, then click **OK**. The completion of one activity starts the execution of the next activity based on the link condition chosen here. If, for example, you connect a **Wait until a situation is True** activity to a **Take action or Write message** activity, and choose **Situation is true** as the condition, the **Take action or Write message** activity begins executing when the situation becomes true. Every activity type except **Suspend execution** gives you multiple conditions from which to choose. You can create an activity for each one of these conditions and connect to each of them from the source activity.

   c. Connect the remaining activities.

      • You can double-click ⊢═⊣ **Connect** to keep it on until you select another activity or click ▷ **Select**. After drawing a connector line, a thin black line is displayed between the two activities. The line disappears when you draw the next connector line or turn off the tool.

      • While ⊢═⊣ Connect is active, every mouse movement in the Grapher view draws the connector line. Right-click if you want to stop drawing temporarily. As soon as you click an activity in the Grapher view again,

the tool will be reactivated and you can continue to draw. Also, you can click ↺ **Undo** to cancel the last change you made (or changes, if you continue to click).

7. When you have finished adding activities and connecting them, click ☑ **Validate** to check the logic flow. If you get an error message, you must fix the problem before you can save the policy.

8. If you have not yet distributed the policy to the managed systems where it will run, select ☑ **Distributed**, select the managed systems or managed system groups to which you want to distribute the policy.

9. When you are finished creating the policy, click **Apply**. Policies are stored at and controlled by the monitoring server to which the managed system is connected, either a remote or the hub Tivoli Enterprise Monitoring Server. The policy is saved and started automatically on the managed systems to which it is distributed unless you cleared the ☐ **Auto start** check box.

10. If the policy is not set to Auto start, start the policy manually by clicking ▶ **Start Policy**.

11. If you are prompted to specify the TEMS to start the policy on, select the monitoring server (or servers) and click **Start**.

Results

The policy is stored at and controlled by the Tivoli Enterprise Monitoring Server to which the agent is connected, either a remote or the hub monitoring server. If the policy distribution is changed manually or dynamically, the policy is restarted. Dynamic updates happen after a new agent comes online or after an agent connection switches from one monitoring server to another.

The first activity in the policy is ☑ **Wait until a situation is True**, which might give the impression that the policy is not running because it is an activity that waits and does nothing until the situation is true. But after that happens, the workflow progresses and continues until all activities have been completed or until the policy is manually stopped with ● **Stop Policy**. If you enabled ☑ **Restart** and the situation is still true, the policy begins again.

If there are invalid activities in the logical workflow, visual feedback is provided showing you which activities are valid and which are not valid. The label of all valid activities changes to blue following validation. If an activity is not valid, the appearance changes in the following ways:

- Border changes from magenta to black
- Label changes from black to red
- Shape changes from rounded rectangle to rectangle
- Tool tip changes to the last error that applies to the activity

Related reference:

"Policy Details" on page 335
Review the Policy Details area of the Workflows window to see the entire list of policies in your monitored environment and information about each policy, such as whether it is distributed or auto started.

"Workflow Editor" on page 338
Use the Workflow editor portion of the Workflows window to create and edit policy workflows by selecting, configuring, arranging, and connecting activities, and to test their validity.

Editing a policy

Use the Policy Details area of the Workflows window for changing policy settings and the Workflow editor area for editing policy workflows.

Before you begin

Your user ID must have Modify Policy permission to use this function.

About this task

To edit a policy:

Procedure

1. If the Workflows window is not open, click ⚙ **Workflow Editor**.
2. In **Policy Details**, click the row whose policy you want to change, then edit the cells as needed. As soon as you make a change, ↩ **Undo** is enabled. Click it if you want to undo the changes made so far to the policy settings.

   After a policy has been saved, you cannot change its name. Instead, copy the policy and edit the name of the copy.
3. Click 📝 **Edit Workflow** to open the policy in the Workflow editor, then edit the activities as needed:
   * To edit an activity, double-click it.
   * To delete an activity or connector, select it and click ✖ **Delete**. You can also use Shift+click to select multiple activities or use click and drag to select a rectangular area.
   * To copy an activity, select it, click ▭ **Copy**. Click where you want it to go in this or another policy open in the work area, and click 📋 **Paste**. You can also use Shift+click to select multiple activities or use click and drag to select a rectangular area.
   * To move an activity, select it, click ✂ **Cut**, then click to where you want to move it in this or another policy open in the work area, and click 📋 **Paste**. You can also use Shift+click to select multiple activities or use click and drag to select a rectangular area.
   * To add an activity, click its icon in the Workflow components frame, then click inside the Grapher view.
   * To connect an activity, click ➡+ **Connect**, then drag from the source activity to the target and click.
   * To restore the Policy Details area, click ▾ Expand; to collapse it, click ▴ Collapse.
4. When you have finished editing, if ☑ **Validate** is enabled, click it to check the logic flow. If there are errors, fix them as indicated in the message, and click ☑ **Validate** again.

5. Click **Apply** to save your changes and continue working with policies; -OR- click **OK** to save your changes and close the window. If the policy is set to ☑ Auto start, it will start running automatically; if not, you will need to click ▶ **Start Policy**.

Copying a policy

Copy a policy when you want to use part or all of an existing workflow in a new policy that you can edit for different activities or settings. Also, use this method when you want to rename a policy.

Before you begin

Your user ID must have Modify Policy permission to use this function.

About this task

To copy a policy:

Procedure

1. If the Workflows window is not open, click 🖧 **Workflow Editor**.
2. In **Policy Details**, click the row of the policy you want to copy.
3. Click 🗋 **Copy**. A new row is added to the list and the policy titled, Copy_of_name, where name is the name of the original policy. The details of the original policy, such as the distribution, are replicated to the new policy.
4. Click inside the **Policy name** field and change the name. After the policy is saved, you will not be able to rename it.
5. Change any other settings for the new policy in the Policy Details area.
6. In the **Workflow editor**, modify the activities as needed:
   - To open an activity for editing, double-click it.
   - To delete an activity or connector, select it and click ❌ **Delete**. You can also use Shift+click to select multiple activities or use click and drag to select a rectangular area.
   - To copy an activity, select it, click 🗋 **Copy**, then click where you want to copy it in this or another policy open in the work area, and click 📋 **Paste**. You can also use Shift+click to select multiple activities or use click and drag to select a rectangular area.
   - To move an activity, select it, click ✂ **Cut**, then click where you want to move it in this or another policy open in the work area, and click 📋 **Paste**. You can also use Shift+click to select multiple activities or use click and drag to select a rectangular area.
   - To add an activity, click its icon in the Workflow components frame, then click inside the Grapher view.
   - To connect an activity, click ⊷+ Connector, then drag from the source activity to the target and click.
7. To save your changes, click **Apply** to keep the window open or **OK** to close it.

What to do next

If you used these steps to rename a policy, ❌ **Delete** the original policy.

Deleting a policy

You can delete a policy that is no longer needed.

Before you begin

Your user ID must have Modify Policy permission to use this function.

About this task

Complete these steps to delete a policy.

Procedure

1. If the Workflows editor is not open, click ⚙ **Workflow Editor**.
2. Select the row that lists the policy you want to delete. You can also select multiple policies with Ctrl+click, or Shift+click to select all applications between the first selection and this one.
3. Click ✖ **Delete Policy**.

Results

After you click **OK** in response to the confirmation message, the policy is deleted from the list. If it was running at the time, it stops; if it was in the middle of executing an activity, that activity will complete.

Starting and stopping policies

You can start or stop a policy as you find necessary. When you create or edit a policy that is set to start automatically, it starts running as soon as you click **Apply** or **OK** in the Workflow editor. If it is not set to Auto start, you can manually start it. You can also stop the policy and restart it manually.

Before you begin

For you to stop or start a policy, your user ID must have Start/Stop Policy permission.

About this task

A workflow can involve many activities and can be in various states for different monitoring agents. Stopping or restarting a policy causes those states to be lost.

You can choose to start or stop a policy on all monitoring agents to which it is distributed or, if your configuration includes remote Tivoli Enterprise Monitoring Servers, on all agents that report to a particular monitoring server. With this added control, you can configure your enterprise in such a way that automation can be stopped for agents that, say, are monitoring subsystems that are undergoing maintenance.

Procedure

1. Open the policy list:
   - In the Navigator Physical view, right-click the ⊞ Enterprise Navigator item or an ◉ Agent Navigator item and click 🗐 Manage Policies from the popup menu. With this method, you can see and select from all the policies that are distributed to the selected managed system.
   - Click ⚙ **Workflow Editor**. With this method, you can see and select all policies, distributed or not distributed.
2. Select the policy that you want to start or stop. You can use Ctrl+click to select additional policies or Shift+click to select all policies between and including the

first selection and this one. In the Manage Policies at Managed System window, you can also click **Select All** for all the policies in the list.

3. Click ▶ **Start Policy** or ⏺ **Stop Policy**.

4. When a message asks you to specify the TEMS list on which to start or stop the policy, select the monitoring server (or servers) and click **Start** or **Stop**.

5. If you are in the Manage Policies at Managed System window, click ⌗ **Refresh** to refresh the view and see your changes.

Results

When you start a policy, it starts on the managed systems to which it is distributed on the selected monitoring server (or servers). All the situations within the policy are also started or restarted.

When you stop a policy, if it was in the middle of executing an activity, that activity will complete. Policies are stored at and controlled by the monitoring server to which the managed system connects. The policy will remain stopped until you start it again manually or, if it is set to **Auto Start**, until you next edit the policy or restart the monitoring server.

If the policy distribution is changed manually or dynamically, the policy is restarted. Dynamic updates happen after a new agent comes online or after an agent connection switches from one monitoring server to another.

Policy distribution

When you create a new policy, it is not initially distributed to any managed systems. This window opens after you select ☑ **Distribution** for a policy in the **Policy Details** list, so you can select the managed systems where you want the policy to run.

If the system is selected and has multiple agents of this type, then distribution will be to all of them. For example, a WebSphere MQ message queue policy written for a system with three WebSphere MQ message queues would be distributed to the three agents of this type on the system.

Assigned
Shows the agents to which the policy is assigned.

Available Managed Systems
Shows the agents of the same type to which the policy can be distributed, such as all UNIX agents for a UNIX situation. To add a managed system to the distribution, click the name, then ◀ to move it to the Assigned list.

Available Managed System Groups
Your product comes with one managed system group named for the product, such as *MS_SQL_SERVER for all MS SQL Server managed systems. To add a managed system group to the distribution, click the name, then ◀ to move it to the Assigned list.

Note: If **Correlate by** is set to Business Application, no individual managed systems are shown; only managed system groups.

Workflows

The Workflows window opens when you click ⚙ **Workflow Editor**. Use this window to see all the policies in your managed environment and to create and edit policy workflows.

The Workflows window has two parts: Policy Details and Workflow Editor.

Policy Details

Review the Policy Details area of the Workflows window to see the entire list of policies in your monitored environment and information about each policy, such as whether it is distributed or auto started.

Toolbar

The Policy Details lists the policies that are stored on the Hub Tivoli Enterprise Monitoring Server and any remote monitoring servers. It has tools for creating and maintaining policies and an area that shows the current version of each Workflow software component.If none of the tools is available, your user ID does not have Modify Policy permission; if, after you select a policy, the Start and Stop tools both remain dimmed, your user ID does not have Start/Stop Policy permission.

▶ **Start Policy**
> Starts the selected policy (or policies) on the managed systems to which it is distributed.

🔴 **Stop Policy**
> Stops running the selected policy (or policies). If the workflow activities are currently executing, processing stops after the current activity is completed.

📄 **New Policy**
> Adds a new row to the Policy list, partially collapses the Policy Details area to show only the new policy, and expands the Workflow Editor area so you can build the workflow. You can restore the policy list with ▾ and collapse it with ▴ , or drag the border up or down. Initially, the policy is named and described as New_Policy. Click inside the field to change the name. Use the Delete or Backspace key to erase to the right or left of the cursor.
>
> If you prefer to keep the Policies Details area open when you create a new policy, use Alt+ 📄 .

📋 **Copy Policy**
> Copies the selected policy. A new row is added to the list, the Policy Details area collapses, and the Workflow Editor area expands to the full window so you can build the workflow. You can manually restore the policy list with ▾ and collapse it with ▴ . Initially, the policy is given the same name as the original, but prefixed with "Copy_of_". In the Policy Details area, click inside the **Name** field to change the name; use Delete or Backspace to erase to the right or left of the cursor.
>
> To keep the Policies Details area open when you copy the policy, use Alt+ 📋 .

✖ **Delete Policy**
> Deletes the selected policy.

◈ **Refresh List**

> If any policies have been added, edited, or deleted by other users in your Tivoli® managed network, you can refresh the Policies list to show the changes.

Policy list

The policy list is a table containing the names and properties of policies on the hub monitoring server and any remote monitoring servers. The first two columns of the table allow you to edit a policy workflow or to undo any changes you have made to the policy during the current session. The other columns contain the names of the policies and properties that you can specify or edit. These columns include the following:

↩ **Undo**

> Voids your changes to the policy in this row.

✎ **Edit Workflow**

> Opens the policy in the Workflow editor for viewing and editing the workflow components. When you do, the **Policies Details** area partially collapses to show only the policy you selected and the Workflow Editor area expands. If you prefer to keep the Policies Details area open, use Alt+ ✎ when you open the Workflow editor.

> Note that you can manually restore the policy list with ▾ Expand and collapse it with ▴ Collapse, or by dragging the border up or down.

Policy name

> The name of the policy, up to 32 letters and numbers, with no spaces or special characters other than the underscore character (_). Click inside this field and edit the text as needed, using the Backspace and Delete keys to erase text to the left or right.

> You cannot rename the policy after it has been saved with **Apply** or **OK**. Instead, copy the policy and change the name and any other settings for the new policy in the **Policy Details** area.

☐ **Distributed**

> Specifies where the policy is to be distributed and run. When you select this option, the Change Policy Distribution window opens with a list of the available managed systems to which you can assign the policy.

> If **Correlate by** is set to Logical Application Group, no individual managed systems are shown; only managed system groups. You must select one from the **Available managed system groups** list and click ◂ .

☑ **Auto start**

> When selected, starts the policy automatically when it is saved and when the monitoring server is started. Clear this option if you want to start the policy manually.

☐ **Save Results**

> When selected, saves the results data internally. You can save results data only to policies that contain a user choice activity.

▾ **Correlate by**

> **Host Name**
>
>> the default mode, maintains separate paths of execution for each host where one or more different managed systems are running. If the host name for two different managed systems is the same, both

managed systems take the same path through the policy. For each distinct path, the policy waits for data from, or runs programs on, the same host that started that path of execution.

Host Address
works identically to Host Name except it identifies the managed systems by their host address instead of host name. It is typically used only if the policy cannot run when it references the host name.

Logical Application Group
identifies the managed systems for processing by their managed system group name. This mode is used to combine disparate agent types. The policy maintains separate paths of execution for each managed system group to which the policy is distributed.

Uncorrelated
which is not normally used, takes one path through the policy for all managed systems to which the policy is distributed. No attempt is made to correlate data from different managed systems as one activity terminates and subsequent activities are started.

☐ **Limit restarts**
Left cleared, specifies no limitation on the number of times the policy can begin running again. When selected, stops the policy after it has run five times in a 24-hour period.

☑ **Restart**
Starts the policy again after the last activity finishes executing. When cleared, stops the policy after the last activity finishes executing until it is started again manually or, if you have ☑ **Auto start** enabled, the next time the policy is saved or when the monitoring server on which it is running is recycled.

Description
Shows the description for the policy. Click inside this field and edit the text as needed, using the Backspace and Delete keys to erase text to the left or right.

Version
The version of the hub monitoring server.

Modified by
Shows the user ID of the person who most recently edited the policy.

Modified date
Shows the time and date the policy was last edited.

Compatibility levels

This area of the window shows ▶ , which you can click to list the minimum version of the Workflow Editor software components that are compatible:
- Workflow Editor
- Tivoli Enterprise Portal Automation Server
- Tivoli Enterprise Monitoring Server Automation Server

Related tasks:

"Creating a policy" on page 328
Use policies to perform actions, schedule work, and automate manual tasks on one or more managed systems in your enterprise.

Workflow Editor

Use the Workflow editor portion of the Workflows window to create and edit policy workflows by selecting, configuring, arranging, and connecting activities, and to test their validity.

When you select a policy from the **Policy Details** list and click ✎ **Edit Workflow**, a Grapher view pane opens in the Workflow editor so you can see and edit the policy activities. You can open multiple policies in the Workflow editor, each in its own pane. This area of the Workflows window has the following components:

- A Toolbar
- A Workflow Components area, with tabs for General activities, Emitter activities, and Extensions on the left
- A Work area and grapher view area where you add and connect activities to create a workflow on the right

Click ▱ (between the Workflow Editor and **Policy Details**) to expand the Workflow Editor area to the top of the window or ▱ to restore the policy list and return the editor to its former size.

Click ◂ Expand (between the Workflow components and the work Grapher view) to expand the work area to the width of the window or ▸ Restore to restore the Workflow components and return the work area to its previous width.

By right-clicking a situation activity in the Grapher view, you can choose to edit the situation specified in the activity. The situation opens in the Situation editor.

Toolbar

☑ **Validate**

Test the workflow to ensure it is composed logically. The text turns blue for activities that validate successfully.

Any errors are reported in a Workflow Definition Error message. The message supplies the activity ID in brackets, which is displayed in a flyover description when you move the mouse pointer over an activity; followed by the name and the error message. You cannot save the policy until you fix any problems and the workflow passes validation.

✂ **Cut**

Removes the selected activity or activities and place on the clipboard.

▢ **Copy**

Copies the selected activity.

▢ **Paste**

Pastes the contents of the clipboard into the editor

✖ **Clear**

Removes the selected activities from the workflow.

Undo

Voids the last action (or actions if you continue to click).

Redo

Repeats the last action.

Select

Selection mode. This is a toggle tool. Click it again to turn off selection. Use the **Select** tool to drag an activity into position or to select it for removal. You can click and drag from one point to the opposite corner to select a group of activities, or use Shift+click to select multiple activities individually.

Perform node layout

Rearranges, as needed, the activities and their connectors to best fit in the view.

Actual size

Returns the view to normal size. Used after zooming.

Fit to Contents

Scales the view so all activities are visible. This is useful when you have added or removed activities.

Zoom in

Zooms in on the view to see an enlarged version.

Zoom out

Zooms the view out to reduce the size and see more in the work area.

Zoom Box

Expands or contracts the view to fit a selected area.

Pan

Pans the view, which selects the entire policy and enables you to move it around the Grapher view.

Related tasks:

"Creating a policy" on page 328
Use policies to perform actions, schedule work, and automate manual tasks on one or more managed systems in your enterprise.

Grapher View

The grapher view is displayed when you create a workflow or have chosen a workflow to edit. Use it to add, edit, and connect activities.

After a workflow activity has been added to the grapher view, you can move the mouse pointer over the activity to see its identifier. If the Workflow was created in the Workflow Editor, the ID will be the activity name followed by a sequence number, such as "WaitOnSituation1" or "SystemCommand3". Otherwise, the ID will look something like "NT285". The activity identifier is referenced in two places:

- In the ⬚ **Take action or Write message** activity, the **More options** button expands the Action Settings dialog. If multiple data items are returned, you can choose to take action on every item generated by the situation in an activity. You can also choose to take action on the agent that returned data for the situation in an activity The activity ID is selected from a list.

- When the workflow has been tested with ☑ **Validate**, if errors are found, the IDs of the affected activities are referenced in the Workflow Definition Error message.

General activities

The **General activities** tab in the Workflow Editor displays the default activities and product-provided predefined activities.

Wait until a situation is True

Opens a list of the situations distributed on your managed network. When you select the situation and click **OK**, the icon with the name of the situation is displayed in the grapher view. When the policy encounters this object in the workflow, it suspends execution until the situation is true. You then must add an activity to perform when the situation fires.

All policies except those that are uncorrelated must begin with this activity.

End code choices for starting the next activity: *Situation is true; Error*.

Evaluate a situation now

Examine the status of an existing situation. Select a situation from the scroll box that is displayed after you drop the activity into the Graphic view. Click **OK** to confirm your selection. Note that the evaluation interval is based on the monitoring interval of the selected situation. There are some special considerations related to the evaluate situation activity described in *Using IBM® Tivoli® Monitoring*.

End code choices for starting the next activity: *Situation is true; Situation is false; Error*.

Take action or Write message

Enables you to issue a system command from a policy or generate a universal message that can be seen in the Universal Message console view.

When issuing a system command, keep in mind that there is a timeout period of 50 seconds. If the command does not complete within that time, an error is issued. As well, the command must be valid on the operating platform on which it will execute.

You can send Simple Object Access Protocol (SOAP) requests. See the *IBM Tivoli Monitoring Administrator's Guide* for details. If the action command is a SOAP request, the policy must be distributed to the hub_ hostname managed system (or *HUB managed system group).

End code choices for starting the next activity: *Error; Action succeeded*.

Make a choice
Not editable

This option is displayed when any of the policies in the list was created in the Candle® Management Workstation (predecessor user interface to the Tivoli Enterprise™ Portal) and includes a Make a choice activity. However, it is not available for adding to policies.

Suspend execution

Specifies a unit of time, expressed in seconds, that you want a policy to wait prior to running. This is a useful feature if, for example, you issued a command, and you want to give the system administrator time to resolve the problem prior to continuing to run the policy.

Start/Stop a policy

Use this activity in a policy to stop another policy from running while the current policy is executing. You can also use to start another policy while the current policy is executing.

End code choices for starting the next activity: *Policy started; Error*.

Start/Stop a situation

Stops a situation from running while the current policy is executing. You can also use this option to start another situation while the current policy is executing.

End code choices for starting the next activity: *Situation started; Error*.

Connector

Adds connecting lines from one activity to the next to control the workflow of the policy.

You can have a single connector from an activity, and you can have multiple connectors from an activity so long as each connector links to a separate activity. Multiple connectors from an activity are processed using AND logic.

When you join activities using the Connector tool, you assign an end code to each connection. As you click and drag from one activity to the next one in succession, a window opens from which you select a link condition. The possible conditions (end codes) resulting from the source activity that will invoke the target activity are listed in the window. For example, if you are connecting a **Wait until a situation is True** activity and choose **Situation is true** as the link condition, the target activity will begin executing when the situation becomes true.

Available end codes include the following:

Action succeeded

The action completed without errors.

Error The activity failed to complete successfully. The reason might be that a referenced object does not exist or a command returned an error status from the operating system.

Resume

The **Suspend execution** activity waits a specified period of time, then resumes executing the policy.

Situation is false

The **Evaluate situation now** and **Wait until a situation is False** activities return this end code if the situation is not currently true.

Situation is true

The **Wait until a situation is True** and **Evaluate situation now** activities return this end code if the situation is currently true.

Choice

The **Make a choice** activity uses this end code to indicate that the linked activity is a potential choice offered to the user when a work item is opened.

Timeout

The **Make a choice** activity uses this end code to indicate the next activity is to run if the user choice times out before a user selects an option.

Action Succeeded

The **Take action or Write message** activity uses this endcode to indicate the action completed without errors.

Action Settings:

In the Workflow editor, clicking ⬚ **Take action or Write message** opens the Action Settings window. Use this function to send a command or universal message to the system. For example, you might want to log information, trigger an audible beep, or stop a job that is using too many resources.

Action Selection
◦ **Universal Message**, when selected, opens this area:

Category of the message you are defining. Type one of the following: critical, warning, or informational.

Severity that you want assigned to the message. Type a severity number (1, 2, and so forth).

Message that is displayed. Type the message text you want to be displayed at the system specified (see Where should the action be run, below).

Universal messages are shown in the ▤ **Universal Message Console** view.

◉ **System Command**, when selected, opens a text box for entering the command to issue at the system. The command language is that of the application and operating platform you are running (such as a UNIX shell command to run at a UNIX system).

Example: Windows has a command for sending a message to a system: net send computer_name message, as in net send MYPC Disk is 85% full. Delete some files.

Example: UNIX has the mail command for directing messages: mail node userID text.

Attribute
A window is opened for choosing an attribute to include in the universal message or system command. In the Select Attribute window, specify an attribute group and item, then click **OK**.

The attribute you selected is displayed in the **Message** field or the **System Command** field. The attribute name is replaced by the actual value during execution of the action. The attributes available for you to select are based on the situations you have added to the policy preceding the action. If there are no situations in the policy, the **Attribute** button is disabled.

In the example that follows, the command sends an email that reads, "System: MyServer, Disk: D:, Space Available: 15%":

```
echo System: &NT_Logical_Disk.Server_Name;, Disk:
&NT_Logical_Disk.Disk_Name;, Space Available:
&NT_Logical_Disk.;%_Free%|mail userID.
```

More Options

If multiple data items are returned
These options determine what action to take if more than one monitored item of data is returned by the situation evaluation:

◉ **Only take action on first item** issues the universal message or the system command on only the first item of data that meets the situation criteria.

◦ **Take action on each item from activity** issues the universal message or system command on all items of data generated by the situation specified in the activity selected from the ▾ list. The ID

for an activity is shown in the Grapher View when you move the mouse pointer over the activity graphic.

The results data returned from each of these individual situation-based activities might originate from the same attribute group on an agent or from attribute groups on different agents. For example, the workflow might be triggered by an OMEGAMON XE for CICS on z/OS situation that monitors for delayed transactions on an CICS monitoring agent, (activity ID *WaitOnSituation1*) and might then evaluate the state of a situation that monitors threads on an OMEGAMON XE for DB2® monitoring agent, (activity ID *EvaluateSituation2*).

Some situations can return a single row and others can return multiple rows for each sample. You can control the number of universal messages or system commands to issue when there are multiple situation rows: a message or a command can be sent for the first row only or for every row. As mentioned earlier, though, there can be multiple sets of situation data originating from different activities. Select this radio button to control which set of results data to use.

Referring to the earlier example, send one command for every row returned by the CICS situation by selecting *WaitOnSituation1*; or send one command for every row returned by the DB2 situation by selecting *EvaluateSituation2*.

☐ **Continue processing rows on failure**, when selected, keeps track of both the rows for which the command failed and those for which it succeeded. A policy can be constructed so that one path is taken for the rows that succeeded and another for the rows that failed.

This is a useful feature when, for example, you want to run a script to filter results. In this case, a failed condition does not mean that the activity itself encountered a problem but that a particular row did not satisfy some criteria. When this option is cleared, a system command that returns a non-zero result while processing a series of situation result rows stops further processing of any remaining rows and the activity is terminated with `Error`, even though some of the commands might have succeeded.

Note: If the situation has a Display Item set and the "Only take action on first item" option is selected, the action is taken once for each unique display item in the rows of data. For example, if there were 100 total rows that qualified and there were 6 unique display items, the situation would fire 6 times. See "Formula" on page 259 for more information on display items.

Where should the action be executed
Determines where to perform the action: on the system where the agent resides, the monitoring server to which it is connected, or the system selected from the list.

◉ **Execute the action at the agent that returned data for activity** issues the universal message or system command on the agent that returned data for the situation in the specified activity (selected from the ▾ list). The ID for an activity is shown in the Grapher View when you move the mouse pointer over the activity graphic.

Referring to the **Take action on each item from activity** example, if the action must be executed at the DB2 agent, you would specify *EvaluateSituation2* because this is the ID of the activity that evaluates the state of the DB2 situation on the DB2 agent.

   ◉ **Execute the action at the Tivoli Enterprise Monitoring Server** issues the universal message or system command at the monitoring server that the agent connects to.

   ◉ **Execute the action at** is used to select a specific managed system from the ▾ list.

Related tasks:
"Adding a Universal Message Console view" on page 145
Create a Universal Message Console view to see the messages produced at the hub Tivoli Enterprise Monitoring Server.

Attribute selection:

The Attribute Selection window lists the available attributes for the group, with those selected listed on the left. This window opens automatically when you select **Attributes** from one of the other panes in the Workflow editor.

Selected Attributes
   Displays the attributes selected for this activity. To remove an attribute from this list, select the attribute you want to remove and click **Remove**. After selecting the first attribute, you can use Ctrl+click to select others or Shift+click to select all attributes between the first selection and this one.

   Remove All selects and moves all the attributes back to the Available Attributes list.

Available Attributes
   Shows the selected attribute group, followed by the attributes available for that group. Click ▾ to see and select from the available attribute groups. The common attributes for Local Time are always available. The other attribute groups are those associated with the emitter type.

   Select the attribute you want from the list and click **Add**. After selecting the first attribute, you can use Ctrl+click to select others or Shift+click to select all attributes between the first selection and this one.

   Add All selects all the attributes in the list and moves them to the **Selected Attributes** list.

Description
   Displays a short description of a selected attribute.

Emitter activities
The **Emitter activities** in the Workflow Editor displays activities associated with alert emitters.

When your managed network includes an IBM Tivoli Alert Adapter or IBM Tivoli Alert Manager, this tab shows the emitter activities available for your products. Emitter activities are used to forward situation events to Tivoli management products and to third-party party management products.

For more information on the parameters for the emitter activity, click **Help** in the Emitter Settings window and refer to the documentation of the target management

product. The SNMP Event activity is available to all users; many of the other activities are available only if your environment has policies that were created with one of the other alert manager products.

NetView® Event

> IBM Tivoli Alert Adapter for TME 10 Netview.

SNMP Event

> This activity enables you to send an SNMP trap with information about any situation event to an SNMP manager.

Connector

> Use the Connector tool to add connecting lines between the activity icons to control the workflow of the policy.

> You can have a single connector from an activity, and you can have multiple connectors from an activity so long as each connector links to a separate activity. Multiple connectors from an activity are processed using AND logic.

> When joining activities with the Connector tool, you assign an end code to each connection. Emitter activities can generate the following end codes:

> - **Error** means that the emitter activity failed to complete successfully. The reason might be that a referenced object does not exist or a command returned an error status from the operating system.
> - **Data emitted** means that the data generated by the activity was emitted successfully to the application.

Emitter settings:

The Emitter activities tab of the Workflow editor contains the available emitter activities for your installed products. After you drop an emitter activity into the grapher view, this window is displayed for you to configure the activity.

Emitter Type
> Shows the name of the emitter activity.

Parameters
> Lists the parameters required for this emitter. Select a parameter to see its current value on the right.

> **Emitter target** identifies the emitter that transfers the request (such as a command or log entry) to the third-party software. The host name is followed by the emitter name, separated by a colon (:), such as myhost:TME_Emit.

> **Attributes** opens the Attribute Selection window, which shows the attributes currently selected to be included in the communication. After viewing and adjusting the list as needed, click **OK** to close the window and see the selected attributes in this area.

> Any other required parameters for your IBM Tivoli Monitoring product will display here. See your product documentation for details.

Current value
> Shows and enables you to change the value or values of the selected parameter.

Operation Mode
> Determines how the communication is transferred to the emitter.

◉ **Invoke emitter once for each data row**, the default, repeats the emitter activity for every row in the data sample that renders the situation true.

○ **Bundle data rows and invoke emitter once** assembles the emitter communication for every row returned in the true situation and sends them as one emitter activity.

For example, consider a situation that fires when the disk space is below 20%. Of the four rows in the data sampling, one for each disk, three meet the condition. The emitter activity sends a message to the third-party event console, depending on the operation mode, once for each of the three disks below 20% or once for all three disks.

Extensions

The **Extensions** tab displays the Wait until a situation is False activity, On Demand Report activity, Send an e-mail activity, and activities associated with Tivoli Application Dependency Discovery Manager (TADDM) integration.

Wait until a situation is False

Waits for the state of a specified situation to change from true back to false, which causes the open event to close. The event closes automatically when the agent is next sampled (determined by the situation sampling interval) and the condition is no longer true.

Wait until a situation is False is useful for policies you want to repeat but only when the state returns to false. An example would be a repeating policy that uses an emitter activity to generate a trouble ticket when CPU activity is high. If this condition continues beyond the 30-minute sampling interval, a new trouble ticket would be generated every 30 minutes. To prevent this from happening the user adds a **Wait until a situation is False** activity to the end of the workflow, so the policy will not restart until the situation has gone to false, then true again.

End code choices for starting the next activity: *Situation is false; Error*.

Note: This activity does not apply to manual event closures.

Configure Tivoli Application Dependency Discovery Manager Integration

This activity is the first step in establishing communication with a TADDM server. There must be one Configuration Tivoli Application Dependency Manager Integration activity for each TADDM server you want to update.

End code choices: *Success; Error*.

See "Configure Tivoli Application Dependency Discovery Manager Integration" on page 347.

Send a Tivoli Application Dependency Discovery Manager Update

Sends information in response to a triggered situation to a TADDM server (previously specified in a **Configure Tivoli Application Dependency Discovery Manager Integration** activity). Use this activity to notify TADDM of resource changes in your monitoring environment. The requests sent can include returned situation data or information from On Demand Report activities.

End code choices: *Success; Error*.

See "Send a Tivoli Application Dependency Discovery Manager Update" on page 349.

On Demand report

Sends an on-demand request to retrieve attribute data for a specified monitoring agent report. You can specify filters to be used in the request, limiting data to the situation that triggered the workflow.

End code choices: *Success; Error*.

See "On Demand Report" on page 350.

Send an e-mail

Sends an e-mail to a single recipient in response to a triggered situation event or activity in your workflow.

End code choices: *Success; Error*.

See "Send an email" on page 352.

Connector

Adds connecting lines from one activity to the next to control the workflow of the policy.

You can have a single connector from an activity, and you can have multiple connectors from an activity so long as each connector links to a separate activity. Multiple connectors to an activity are processed using AND logic.

When you join activities using the Connector tool, you assign an end code to each connection. As you click and drag from one activity to the next one in succession, a window opens from which you select a link condition. The possible conditions (end codes) resulting from the source activity that will invoke the target activity are listed in the window. For example, if you are connecting a **Wait until a situation is True** activity and choose **Situation is true** as the link condition, the target activity will begin executing when the situation becomes true.

Related concepts:

"Tivoli Application Dependency Discovery Manager policies" on page 353
You can create policies that interact with your Tivoli Application Dependency Discovery Manager (TADDM) environment. These policies can forward updates such as, new, changed, or destroyed resources to your TADDM server.

Configure Tivoli Application Dependency Discovery Manager Integration:

Selecting the **Configure Tivoli Application Dependency Discovery Manager Integration** activity opens a dialog box for entering the information required to send updates to the TADDM server.

Define how Tivoli Monitoring communicates with TADDM and prepares TADDM to accept subsequent update requests from your Tivoli Monitoring environment.

If you have multiple TADDM servers, you must provide information for each server using separate **Configure Tivoli Application Dependency Discovery Manager Integration** activities. Include this configuration activity in the workflow of a policy that is defined with **Auto start** ☑ selected, **Correlated by** set to **Uncorrelated**, and **Restart** ☐ *not* selected.

If you change any fields at a later date, the new values become effective immediately (assuming the containing policy has been defined to **Auto start**); the next update sent by any policy to this TADDM server uses the new values.

Tivoli Application Dependency Discovery Manager Server Alias

The server alias represents a TADDM server. The server alias specified here must be used by all activities that intend to send updates to this TADDM server.

For example, the server alias specified for the associated **Send Tivoli Application Dependency Discovery Manager Update** activities must match the server alias specified here.

The server alias can be an arbitrary name such as, `Prod1`, `Test2`, or `Austin3`.

Tivoli Application Dependency Discovery Manager Server Type

Select the corresponding server type of the TADDM server to which you are sending updates. The server type depends on how the TADDM server is configured. Select one of the following options: Synchronization or Domain/Storage. See Three ways of deploying TADDM in the TADDM Information Center for more details.

Management Software System Hostname

This hostname must match the hostname provided by the Discovery Library Adapter (DLA) discovery process. This value is used to construct the MSSName that is used in subsequent TADDM update activities. The MSSName is used by the TADDM server to correlate update notifications with information provided through the discovery process.

Typically, the value should specify the hostname of the portal server found in the DLA XML output file. For example, if the DLA XML output file includes `<cdm:Hostname>aix1.lab.raleigh.ibm.com</cdm:Hostname>`, then `aix1.lab.raleigh.ibm.com` would be the hostname you specify.

For more information on the DLA process see "Using the Tivoli Management Services Discovery Library Adapter" in the *IBM Tivoli Monitoring Administrator's Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/adminuse/itm_admin.htm).

Tivoli Application Dependency Discovery Manager Server URL

Enter the information for the TADDM server. The protocol options are http (default port 9430) or https (default port 9431). You can also specify your own valid port number.

Tivoli Application Dependency Discovery Manager User Credentials

Required user name and password for the TADDM server.

Tivoli Enterprise Portal Server URL

Enter the information for the Tivoli Enterprise Portal Server. This URL is stored in the TADDM server and used to perform any launch-in-context activity. The protocol options are http (default port 1920) or https (default port 3660). You can also specify your own valid port number. Specify the same hostname and port number that your portal server is configured to use.

Request Parameters

The Request Parameters section is used to define the format of the JavaScript Object Notation (JSON) update request that IBM Tivoli Monitoring sends to the TADDM server. These three fields are included: Key, Value Type, and Value.

This section is predefined and typically does not require modification.

If needed, make a copy of the TADDM_Initialization policy and use the request parameters included in the Configure Tivoli Application Dependency Discovery Manager Integration activity.

Insert Variable Data

Because the Request Parameters section is typically predefined, you do not need to use substitution options.

Request timeout

Select the amount of time in seconds you want to allow before a timeout occurs. A timeout occurs when IBM Tivoli Monitoring doesn't receive a response from the TADDM server. Options (in seconds): 5, 10, 15, 30 (default), 45, 60, 90, 120.

Send a Tivoli Application Dependency Discovery Manager Update:

Selecting the **Send a Tivoli Application Dependency Discovery Manager Update** activity opens a dialog box for entering information required to forward changes in your monitoring environment, discovered through situation events, to a TADDM server.

Tivoli Application Dependency Discovery Manager Alias

The server alias represents a TADDM server. Use the same server alias that you specified in the associated **Configure Tivoli Application Dependency Discovery Manager Integration** activity. The server alias can be an arbitrary name such as `Prod1`, `Test2`, or `Austin3`.

Operation Type

The operation type determines what type of update is sent to the TADDM server. Select one of the following types appropriate to the situation that triggered the workflow: Create, Update, Move, or Delete.

Request Parameters

The Request Parameters section is used to define a JavaScript Object Notation (JSON) update request that IBM Tivoli Monitoring sends to the TADDM server. These three fields are included: Key, Value Type, and Value. Monitoring agents that exploit this activity provide policies that include predefined request parameters.

The Key contains a list of common predefined keys. Other valid key values can be entered, but they must be part of the Tivoli Common Data Model (defined in the IBM Multilingual Terminology database).

The Value Type is a fixed list of selectable values.

The Value can be a predefined selected value or a user-defined value. The most common options are displayed.

See "Request Parameters" on page 354 for more information.

Insert Variable Data

Specify **Situation Attribute** or **Other value** to substitute data from a preceding situation or On Demand activity.

The **Situation Attribute** button becomes active if the policy contains a "Wait until a situation is true" activity and you select a row cell under Value.

The **Other value** button becomes active if the user selects a row cell under Value. Selecting this button opens a dialog box that facilitates the substitution of data returned by a preceding On Demand activity, as well as diagnostic information regarding TADDM activities and On Demand activities.

Request timeout

Select the amount of time in seconds you want to allow before a timeout

occurs. A timeout occurs when IBM Tivoli Monitoring doesn't receive a response from the TADDM server. Options (in seconds): 5, 10, 15 (default), 30, 45, 60, 90, 120.

If multiple data items are returned
Your selection here controls how this activity behaves if the triggering situation returns multiple rows of data. The selection should be predefined based on the supported monitoring agent.

If an update request should be sent only for the first situation row, select **Execute request once**. Otherwise, select **Execute request for each row from activity** and select either the situation that triggered the workflow or the On Demand activity. Ensure the **Continue processing rows on failure** option is selected.

On Demand Report:

Use the **On Demand Report** activity to request real-time data from any monitoring agent. This activity is generally used to retrieve additional information related to the situation that triggered a workflow.

ITM SOAP Server URL
The default value is **Localhost**. **Protocol** options are http (default) or https. Select **Localhost** or **Other**. When selecting **Other**, enter the specific hub monitoring server where the targeted SOAP server is running and to which the On Demand Report will be sent. You must also specify the valid port number used by the SOAP server.

Application
Select the application. For example, `Windows OS`.

Attribute Group
Select the attribute group. Your options depend on the application you previously selected.

Data source
For Automatic, the On Demand Report request is sent to the monitoring agent that triggered the workflow. For Other, you can specify a valid managed system name to which the On Demand Report request will be sent.

ITM SOAP Server User Credentials
Enter the user name and password to access the ITM SOAP server. Credentials are required if authentication is enabled at the hub monitoring server where the targeted SOAP server is running.

Attributes
Select the attributes you want to be returned in the report. You must select as least one attribute.

Filter Add optional filters to the report. The report returns only rows that meet the filter criteria.

The **Attribute** element is a list of attributes based on the selected **Attribute Group**.

The supported operator values are Equal to, Not equal to, Less than, Less than or equal to, Greater than or equal to, Greater than, and Like. For the Like operator, the '*' character matches one or more characters. Only supported for text attributes.

The **Value** element can be a hard-coded literal or the substitution of an attribute or value returned by a preceding Wait On Situation or On Demand activity.

Multiple filters are supported as conjuncts, for example, using AND to join together.

Insert Variable Data

Specify **Situation Attribute** or **Other value** to substitute data from a preceding situation or On Demand activity.

The **Situation Attribute** button becomes active if the policy contains a "Wait until a situation is true" activity and you select a row cell under Value.

The **Other value** button becomes active if the user selects a row cell under Value. Selecting this button opens a dialog box that facilitates the substitution of data returned by a preceding On Demand activity, as well as diagnostic information regarding TADDM activities and On Demand activities.

Request timeout

Select the amount of time in seconds you want to allow before a timeout occurs. A timeout occurs when IBM Tivoli Monitoring doesn't receive a response from the SOAP server. Options (in seconds): 5, 10, 15 (default), 30, 45, 60, 90, 120.

If multiple data items are returned

Your selection here controls how this activity behaves if the triggering situation returns multiple rows of data. The selection should be predefined based on the supported monitoring agent.

If an update request should be sent only for the first situation row, select **Execute request once**. Otherwise, select **Execute request for each row from activity** and select either the situation that triggered the workflow or the On Demand activity. Ensure the **Continue processing rows on failure** option is selected.

Value Substitution:

When you are using **Other value** data substitution, the following options are available.

The default substitution values are for an On Demand activity.

Substitution type

Your selection here dynamically changes the window options.

Value: Use this type to substitute data returned by an On Demand activity.

Diagnosis: Use this type to substitute diagnosis data.

Other: Use this type to substitute other values from the active policy.

Web Services Activity

Displays when **Value** and **Diagnosis** are selected as the substitution type. Select the appropriate activity from the drop-down list.

Value Type

Displays only when **Value** is selected as the substitution type. Your selection here dynamically changes the window options.

On Demand attribute: Use this type to substitute an attribute from an On Demand activity. In the **Attribute name** field, specify the name of the attribute returned by the On Demand activity.

> **Note:** Leave out any non-alphanumeric characters and replace all spaces with underscores (_). If an attribute name begins with a non-alphanumeric character followed by a space, omit both. For example, `% User Time` is entered as `User_Time` and `Avg % Processor Time` is entered as `Avg__Processor_Time`.

XML element: Use this type to substitute XML data. Fill out the now visible **Element substitution** section providing the appropriate attribute and XML data. Typically this is not used.

Diagnosis type

Displays when **Diagnosis** is selected as the substitution type. Select one of the following options:

Request: Use when you want the content of the request sent by the specified Web Services Activity.

Response: Use when you want the response returned by the request. No response is available for the **Configure Tivoli Application Dependency Discovery Manager Integration** activity or the **Send a Tivoli Application Dependency Discovery Manager Update** activity; N/A is returned if these activities are selected as the web service.

Status: Use when you want the HTTP status of the request.

Other values

Displays only when **Other** is selected as the substitution type. Select the appropriate value from the drop-down list.

Send an email:

Use the **Send an email** window to send notification of a situation event or activity in your workflow to a single recipient. This activity does not support multiple email recipients. This activity is not supported on a z/OS monitoring server.

However, if the hub monitoring server is not running on z/OS, a policy running on a z/OS remote monitoring server can be configured to send emails through the hub's SOAP server.

ITM Web Server URL

The default value is **Localhost**. **Protocol** options: http (default) or https. Only http and https are supported. Select **Localhost** or **Other**. When selecting **Other**, enter the specific hub monitoring server where the targeted SOAP server is running and to which the email request will be sent. You must also specify the valid port number used by the SOAP server.

SMTP Server

Your email server. For example, `smpt.us.ibm.com`.

Sender

The email address that is sending the notification.

To   The email address of a single recipient. Separate a list of recipients with commas.

Subject

Specify the subject heading of the email. This field is optional.

Text Specify any text to be included in the email. This field is optional.

Insert Variable Data

Specify **Situation Attribute** or **Other value** to substitute data from a preceding situation or On Demand activity.

The **Situation Attribute** button becomes active if the policy contains a "Wait until a situation is true" activity and you select a row cell under Value.

The **Other value** button becomes active if the user selects a field. Selecting this button displays a dialogue that facilitates the substitution of data returned by a preceding On Demand activity, as well as diagnostic information regarding TADDM and On Demand activities.

If multiple data items are returned

Your selection here controls how this activity behaves if the triggering situation returns multiple rows of data.

If an update request should be sent only for the first situation row, select **Execute request once**. Otherwise, select **Execute request for each row from activity**, then select the situation that triggered the workflow or the On Demand activity. Ensure the **Continue processing rows on failure** option is selected.

Tivoli Application Dependency Discovery Manager policies

You can create policies that interact with your Tivoli Application Dependency Discovery Manager (TADDM) environment. These policies can forward updates such as, new, changed, or destroyed resources to your TADDM server.

Note: This functionality is only supported for the IBM Tivoli Monitoring for Virtual Environments Agent for VMware VI.

Two types of policies are associated with TADDM:

Initialization

Initialization policies establish communication with a TADDM server and only contain the **Configure Tivoli Application Dependency Discovery Manager Integration** activity. For every TADDM server you wish to update, you must have a matching **Configure Tivoli Application Dependency Discovery Manager Integration** activity.

IBM Tivoli Monitoring provides the TADDM_Initialization policy as a template. This policy includes one **Configure Tivoli Application Dependency Discovery Manager Integration** activity.

To set up your TADDM integration, complete the following steps:

1. Click 🖳 **Workflow Editor**.
2. Select the **TADDM_Initialization** policy, and click 🗐 **Copy Policy**. Give the new policy a meaningful name, such as TADDM_Init_Austin3.
3. In the workflow for your new policy, double-click the **Configure Tivoli Application Dependency Discovery Manager Integration** activity and enter your specific server information. Fill out any required fields (indicated by ???).
4. Ensure the policy is configured as follows:

Policy name	Distributed	Auto start	Save results	Correlate by	Limit restarts	Restart
TADDM_Init_Austin3	✔	✔	☐	Uncorrelated	☐	☐

When selecting **Distributed**, make sure to assign all monitoring servers (select the *ALL_CMS managed system group).

5. To save your changes, click **Apply** or **OK**.

Making a copy of the original TADDM_Initialization policy template ensures that any future updates to this policy template do not overwrite your customization.

Since initialization policies are uncorrelated and not triggered by a situation, they start with the monitoring server and establish a connection to the TADDM server before the update policies start.

If your environment includes multiple TADDM servers, you can group multiple **Configure Tivoli Application Dependency Discovery Manager Integration** activities in one initialization policy. These activities do not need to be connected to each other. Choosing to group your **Configure Tivoli Application Dependency Discovery Manager Integration** activities might be a good idea for administrative maintenance.

The **Configure Tivoli Application Dependency Discovery Manager Integration** activity defines values that are later used by **Send a Tivoli Application Dependency Discovery Manager Update** activities.

Update

Update policies forward changes in your monitoring environment to a TADDM server in response to a triggered situation. Update policies contain the **Send a Tivoli Application Dependency Discovery Manager Update** activity, and might contain other activities such as **On Demand Report** and **Send an e-mail**.

Monitoring agents that support TADDM integration provide predefined update policies. Information for the activities within predefined policies might need to be customized for your environment.

You can have multiple update policies. Each policy might perform a different update operation (for example: Create, Update, Move, or Delete).

Complete the following steps to modify any of the policies:

1. Select the policy, and click 📝 **Edit Workflow**.
2. Double-click any activity in the Grapher View to edit the parameters.

TADDM integration activities are located on the **Extensions** tab in the Workflow Editor.

Note: For z/OS, ICSF (Integrated Cryptographic Service Facility) must be enabled to allow password encryption.

Related reference:

"Extensions" on page 346
The **Extensions** tab displays the Wait until a situation is False activity, On Demand Report activity, Send an e-mail activity, and activities associated with Tivoli Application Dependency Discovery Manager (TADDM) integration.

Request Parameters

The Request Parameters table is used to define a JavaScript Object Notation (JSON) update request that identifies IBM Tivoli Monitoring to the Tivoli Application Dependency Discovery Manager (TADDM) server or that defines the updated resource information that IBM Tivoli Monitoring sends to the TADDM server.

Monitoring agents that support TADDM integration will provide predefined policies containing TADDM activities that include predefined request parameters.

There are three fields: Key, Value Type, and Value. An update request must begin with a start array and end with an end array. An array can contain one or more objects. Each object is started with a start object and ended with an end object. Within an object every value type has a key.

- The **Key** contains a list of common predefined keys in a drop-down menu. Other valid key values can be entered, but they must be valid attributes supported by the associated Tivoli Common Data Model (CDM) classes. To access the CDM documentation see the TADDM Software Developer's Kit (SDK).

- The **Value Type** is a fixed list of selectable values from a drop-down menu.

- The **Value** can be a predefined selected value or a user-defined value. The most common options are displayed in a drop-down menu.

  When providing a value in the Value column for a given Key, you might need to supply the associated IBM Tivoli Monitoring sensor that is provided by TADDM. See the "Sensor reference" topics in the Tivoli Application Dependency Discovery Manager information center.

Examples

Configure Tivoli Application Dependency Discovery Manager Integration:
The predefined TADDM_Initialization policy includes the following JSON request for the Configure Tivoli Application Dependency Discovery Manager Integration activity:

Key	Value Type	Value
	Array	Start
	Object	Start
_cdmclass	String	com.collation.platform.model.topology.process.ManagementSoftwareSystem
_jsonid	String	_jsonid_0
mssname	String	&{=(*TADDM_=Alias+_MSSNAME)}
hostname	String	&{=MSSHostname}
manufacturename	String	IBM
productname	String	IBM Tivoli Monitoring Services
label	String	IBM Tivoli Monitoring
	Object	End
	Object	Start
_cdmclass	String	com.collation.platform.model.topology.process.MSSAccessInfo
_jsonid	String	_jsonid_1
accessString	String	&{=SourceContactInfo}
accessType	String	default
parent	String	_jsonid_0
	Object	End
	Array	End

Send a Tivoli Application Dependency Discovery Manager Update with **Operation Type: Create**
The VMware VI monitoring agent provides a policy to update TADDM when new resources are created. The following JSON update request is predefined for the Send a Tivoli Application Dependency Discovery Manager Update activity:

Key	Value Type	Value
	Array	Start
	Object	Start

Key	Value Type	Value
_cdmclass	String	com.collation.platform.model.topology.sys.vmware.VmwareUnitaryComputerSystem
_jsonid	String	_jsonid_1
uuid	String	&WaitOnSituation1:KVM_TOPOLOGICAL_EVENTS.Host_UUID
childsystem	String	_jsonid_2
	Object	End
	Object	Start
_cdmclass	String	com.collation.platform.model.topology.sys.UnitaryComputerSystem
_jsonid	String	_jsonid_2
virtual	Boolean	true
name	String	&{CT_Get2:/VM_Name}
label	String	&{CT_Get2:/VM_Name}
vmid	String	&{CT_Get2:/MORef}
uuid	String	&WaitOnSituation1:KVM_TOPOLOGICAL_EVENTS.VM_UUID
memorysharesvalue	String	&{CT_Get2:/Memory_Shares}
numcpus	String	&{CT_Get2:/Num_CPUS)
cpusharesvalue	String	&{CPU_Shares)
hostsystem	String	_jsonid_1
	Object	End
	Array	End

Send a Tivoli Application Dependency Discovery Manager Update with
Operation Type: Delete
The VMware VI monitoring agent provides a policy to update TADDM when
resources are deleted. The following JSON update request is predefined for the
Send a Tivoli Application Dependency Discovery Manager Update activity:

Key	Value Type	Value
	Array	Start
	Object	Start
_cdmclass	String	com.collation.platform.model.topology.sys.windows.WindowsComputerSystem
_jsonid	String	_jsonid_2
uuid	String	&WaitOnSituation1:KVM_TOPOLOGICAL_EVENTS.VM_UUID
	Object	End
	Array	End

Send a Tivoli Application Dependency Discovery Manager Update with
Operation Type: Move
The VMware VI monitoring agent provides a policy to update TADDM when
resources are moved. The following JSON update request is predefined for the
Send a Tivoli Application Dependency Discovery Manager Update activity:

Key	Value Type	Value
	Array	Start
	Object	Start
_cdmclass	String	com.collation.platform.model.topology.sys.vmware.VmwareUnitaryComputerSystem
_jsonid	String	_jsonid_1
uuid	String	&WaitOnSituation1:KVM_TOPOLOGICAL_EVENTS.Host_UUID
childsystem	String	_jsonid_2
	Object	End
	Object	Start
_cdmclass	String	com.collation.platform.model.topology.sys.windows.WindowsComputerSystem
_jsonid	String	_jsonid_2
virtual	Boolean	true
uuid	String	&WaitOnSituation1:KVM_TOPOLOGICAL_EVENTS.VM_UUID
hostsystem	String	_jsonid_1
	Object	End
	Array	End

Take action commands

The take action feature enables you to invoke a predefined command or enter a command on a selected managed system. Action commands can also be added to situations to implement simple (reflex) automation when an event is opened.

Some monitored applications provide predefined take action commands. You can customize those commands and create command definitions of your own, then invoke them on the system you choose.

Tip: If you want to send a command to another monitored network, use the Tivoli Enterprise Monitoring Web Services to send Simple Object Access Protocol (SOAP) requests to another environment. See the *IBM Tivoli Monitoring Administrator's Guide*.

Creating a take action command

You can create take action definitions and save them, then invoke them at will through the pop-up menu for Navigator item or query-based view, or automatically when a situation event is opened.

Before you begin

Your user ID must have Modify permission for the Action feature.

Note: The action command is issued with the authorization of the user ID on the monitoring agent or server computer to which it is sent. If your hub Tivoli Enterprise Monitoring Server is set to ☑ **Security: Validate User**, the default, then your ID must be authorized on the relevant system receiving the command. For example, to issue a TSO command, your user ID and password must be a valid TSO ID and it must also be defined as a valid logon ID on the portal server and on the domain or operating system where the hub monitoring server resides.

Tip: You can employ Tivoli Enterprise Monitoring Web Services to send Simple Object Access Protocol (SOAP) requests to other Tivoli Management Services environments. See the *IBM Tivoli Monitoring Administrator's Guide* for details.

About this task

Take action commands can be issued either by typing a command into a text box or by selecting a predefined command from a list. Take action commands can also be added to situations to implement simple (reflex) automation. Some applications provide predefined commands. You can customize those commands or create predefined commands of your own and invoke them as needed on the system you choose. To create or edit a custom Take Action command:

Procedure

1. Select the Navigator item associated with the application on which you want to run the command (for example, Windows OS or z/OS operating system).
2. Right-click one of the following:
   - The selected Navigator item
   - Row in a table view or event console view
   - Data point in a chart
   - Object of a topology view

3. Click **Take Action** → ⚡ **Create or Edit**. The Edit Action and Select Action dialogs open.

4. Select **<Create new Action>** and click **OK**.

5. In the Create New Action window, type a name and description for the command.

6. If there is a choice of Monitored Application, select one from the list.

7. Select the type of command. `System Command` is for issuing a command on the operating system associated with this Navigator item. For example, if you selected a Navigator item on the UNIX Systems branch, you can enter a UNIX shell command.

   The other types of commands that appear are for the monitoring agents associated with this Navigator item. For example, at the Enterprise level, you will see all agent types on your managed network, whereas at the system level, you will see those for the types of agents running on that system.

8. In the text box, type the command just as you would from the command line. The character limit is 512 bytes, including any attribute substitution you specify in the next step.

9. If you want to insert an attribute variable, which will be replaced by its value at run time, click **Insert Attribute**, then select one from the Select Attribute window. Example: This command sends an email that reads, "System: MyServer, Disk: D:, Space Available: 15"

   ```
   echo System: &NT_Logical_Disk.Server_Name, Disk: &NT_Logical_Disk.Disk_Name,
   Space Available: &NT_Logical_Disk.%_Free%|mail userID.
   ```

   The attribute variables will be resolved if the action is launched from a table row or plot point (such as a bar in a bar chart) that can provide the values. If an attribute (or any other string preceded by an **&** ampersand) is put into a take action definition, it will be resolved unless no value is found, in which case you are prompted to enter a value. If none is provided, it will be excluded from the generated output.

10. When you are finished creating the action, click **OK** to save it. Your user-defined action is now selectable from the list of actions available for this agent. The action will be available to all users whose IDs have View Action permission and the application the command was written for as one of their Allowed Applications.

11. Select a name from the list and click **OK**.

12. In the Edit Action window, edit the command name, description, type, or text.

13. If you want to insert an attribute name as a variable that will be replaced by the attribute's current value when the command is invoked, click **Insert Attribute** and select from the list.

14. When you are finished editing the command, click **OK** to save your changes.

Editing a take action command

You can open the Edit Action window to see the list of saved take action commands and then select the one that you would like to modify.

Before you begin

Your user ID must have Modify permission for the Action feature.

About this task

Complete these steps to modify an existing take action command definition.

Procedure

1. Select a Navigator item associated with the application or system on which the command can be run.
2. Right-click the Navigator item or a data point, row, or object in a query-based view.
3. Click **Take Action** > ⚡ **Create or Edit**. The Edit Action and Select Action windows open.
4. Select a name from the list and click **OK**.
5. In the Edit Action window, edit the command name, description, type, or text.
6. If you want to insert an attribute name as a variable that will be replaced by the attribute's current value when the command is invoked, click **Insert Attribute** and select from the list.
7. When you are finished editing the command, click **OK** to save your changes.

Sending a take action command

You can use the take action feature to initiate a command on a managed system in your monitored enterprise. For example, you can restart a process that is not functioning properly or you can send a message.

Before you begin

Your user ID must have View permission for the Action feature.

Your user ID must be identifiable and authorized on the computer where the action will take place.

About this task

To send a take action command, complete these steps:

Procedure

1. Take one of the following actions:
   - Select the Navigator item associated with type of command you want to run or the managed system on which you want to run the command.
   - Open the workspace containing the view from which you want to launch the application.
2. Right-click one of the following:
   - Navigator item.
   - Row in a table or event console view.
   - Data point in a chart (such as a pie chart slice)
   - Topology view object.
3. Click **Take Action** > ⚡ **Select**.
4. Select a predefined action or enter a command using the **Command** field.

   To select a predefined action:
   - a. In the **Name** field click the ▾ **List** button to see and select an action name.

The command is displayed in the field below. Some actions do not require you to add a value. Those that do open the **Edit Argument Values** window automatically. Other actions might display a suggested value, which you can change by clicking **Arguments** and entering a different value.

b. If the **Edit Argument Values** window opens, enter a value to identify the item on which the action must take place.

c. If the action command fills in the value automatically and you want to change it, click **Arguments** and enter the value that identifies what to act on.

A two-column window opens with the attribute in the **Name** column and a place to enter the value of the attribute in the **Value** column. For example, if you chose the action that starts a process, you must identify the process by entering its name as the value. If no argument is required for the action command, this function will be disabled.

d. Click an application host in the **Destination System** list.

If you want to send the command to multiple systems, use Ctrl+click to select more or Shift+click to select all systems between.

To enter a command:

a. In the **Command** text box, type a command. For example:

```
Net send mars The server is running out of memory.
Close some applications
```

The user at the system named mars sees the message, "The server is running out of memory. Close some applications."

b. Click an application host in the **Destination System** list.

If you want to send the command to multiple systems, use Ctrl+click to select more or Shift+click to select all systems between. After you click OK, the command is run on the selected system or systems. If a message is returned, it will be from the application on which the command was run.

Creating a take action view

Tivoli Enterprise Portal enables you to interact directly with your applications and operating system through its take action feature. You can use the take action view to enter a system command or to stop or start a process at the system.

Before you begin

Your user ID must have Workspace Author Mode permission to use this function.

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new pane, click ▥ Split Vertically or ▤ Split Horizontally in one of the views. The pane divides and a duplicate is created. If you split the Navigator, the new view is empty.
3. Click ⚡ **Take Action**. The mouse pointer changes to the view icon (hand icon on Linux). If you get the wrong icon or decide not to add the view, press Escape or click somewhere in the toolbar.
4. Click inside the view where you want the take action view. After you click, the old view is replaced by the take action view.

What to do next

If you want to keep the take action view in this workspace for future work sessions, either click 🔖 **Save** to update the workspace properties with the new view, or click **File** > **Save Workspace As** to save this as a new workspace and leave the original unchanged.

Rules and guidelines for take action commands

Follow these rules when specifying arguments and commands for a take action definition.

- A single parameter should not contain interleaving blanks, unless you enclose the text in double quotes ("). Arguments containing interleaving blanks are separated into individual values. Example: F HCHECKER,UPDATE,CHECK=(&Health_Checker_Checks.Check_Owner, &Health_Checker_Checks.Check_Name),PARM="&Health_Checker_Checks. Parameters"

- Specify substitution variables by prefixing attribute names with an & ampersand.

- Specify fully qualified attributes in the *&Groupname.Attributename* format, indicating that a value can only be substituted from a specific attribute group.

- Command strings are created by concatenating the contents of the Command entry field with all the arguments in the Arguments field.

- When an attribute is selected and added to the Arguments field, the attribute is removed from the list of available attributes.

- A command string (Command + Arguments) definition is limited to 512 bytes. The formatted command string, after variable substitution, is limited to 256 bytes.

- The action command is issued with the authorization of the user ID on the managed system or monitoring server machine to which it is sent. If your managed network validates user IDs, then your ID must be authorized on the relevant system receiving the command. For example, to issue a TSO command, your user ID must be a valid TSO ID and it must also be defined as a valid logon ID for the Tivoli Enterprise Portal.

Launch application

Use the launch application feature to write and save definitions that start applications on your computer. For example, using single sign-on in the browser client or Java Web Start client, you can launch into another participating Tivoli web application.

A launch definition can include arguments that are passed to the application when it is started. The definition can be an operating system command or URL, and can include arguments. A simple example is opening a specific text file in a text editor. And another example is opening your browser to access the RMF™ Monitor III Portal.

You can launch an application by creating a definition right when you need it or by selecting from a list of existing definitions. Some monitoring products provide predefined launch definitions. You can customize those definitions or use the launch application feature to build a set of command definitions for yourself or for all users in your monitored environment.

The Tivoli Enterprise Portal single sign-on (SSO) feature enables you to launch out of the portal into other Tivoli web-based or web-enabled applications without having to re-enter your user credentials. From the portal browser client or Java Web Start client, you can launch into another participating Tivoli web application.

Creating a launch definition

Create a custom launch definition from the Navigator item or query-based view (including event console views) where you want to enter the command. When launched from a query-based view, the command can include attribute values from the launch point, such as IP address or process name.

Before you begin

If your monitored application has predefined launch application definitions or you would like to include attribute names in the argument as symbols to be replaced by the attribute values at launch time, see the product user's guide for the settings.

Your user ID must have View and Modify permission for the Launch Application feature. If you want the new launch definition to be available to other users, you must be in Workspace administration mode when you create it.

About this task

Take these steps to create a launch definition:

Procedure

1. Select the Navigator item to launch from.
2. If the launch definition will originate from a table, chart, event console view, or relational-based topology view, open the workspace containing the view. If the workspace is not the default, use the ⊞ Workspace Gallery to select it. You can add an argument that includes attribute names in the definition that are resolved at launch time with the current attribute values.
3. Right-click one of the following:
   - Navigator item.
   - Row in a table or event console view.
   - Data point in a chart (such as a plot chart point).
   - Topology view object.
4. Click ▶ **Launch** and, in the Create or Edit Launch Definitions window, click **Create New**.
5. In the **Name** text box, type a descriptive name.
6. In the **Target** text box, click **Browse** to locate and select the application to launch. If you leave the field empty, the default behavior is to open Windows Internet Explorer as the launch application. This is desirable when a URL is included in the definition and it will be launched from different computers. If you want a different browser used, the kjr.browser.default client parameter can be changed through Manage Tivoli Enterprise Monitoring Services, as described in the *IBM Tivoli Monitoring Administrator's Guide*.
7. In the **Arguments** text box, type any arguments to be added to start the application:
   - To launch a web page, enter the URL, such as `http://ibm.com`.
   - To add a file name, click **Browse**, select ◉ **Filename**, then locate and select the file.

If any of the argument contains spaces, surround it with double quotes, as in `"C:\Documents and Settings\Administrator\My Documents\db_notes.doc"`. If you have multiple files to specify, separate each entry with a space and enclose in double quotes, as in `"C:temp\my_log.txt"` `"C:temp\your_log.txt"`.

- To add an attribute or other symbolic reference that is known at this location, click **Browse**, select ◉ **Substitutable Items**, then select an attribute from the 📂 **Selected Context** list or a symbol from the 📂 **Extended Context** list. The value will be retrieved at launch time. For example, 👤 **User Id (As Typed)** shows as `&-12002` in the argument and evaluates to `sysadmin`; 📊 **Process Name** shows as `&NT_Process.Process_Name` in the argument and evaluates to `BigProcess`.
- To prompt for a value, type a `^` caret mark then the prompt, as in `^Name`, followed by a single space or a comma. If you would like to supply a default value, but also enable the user to change it, append the prompt with an = equal sign and the value. For example, `^Name=Fred`, where `Name` is the prompt and `Fred` is the default value.
- To prompt for a hidden value, such as a password, type two `^^` caret marks then the prompt (no spaces are allowed), as in `^^Password`.
- To add data or an attribute to look up in context, precede the data with an & ampersand.
- To include an & ampersand or `^` caret in the filename or value, precede it with a / forward slash. For example `john&jane.lwp` would be entered as `john/&jane.lwp`.

8. If you added arguments, you can give the user the option to change the values at launch time by selecting ☑ **Edit arguments before execution**.
9. If you need to specify the **Start in** location, click **Browse** to locate and select the folder for the program or other required files.
10. If the definition originates from a Navigator item and you want to restrict its availability to this item only or to others of the same type, click **Advanced** to open the Advanced Options for Launch Definition window and select one of the options.
11. Click **Evaluate** to see how the arguments will resolve when the application launches.
12. When you are finished with the launch definition, click **Apply** to save your changes and keep the window open; or click **OK** to save your changes and close the window.

Copying or editing a launch definition

You can use an existing launch definition as a template for creating another definition and edit the copy. The predefined definitions for your product can offer a good starting point.

Before you begin

Your user ID must have View and Modify permission for the Launch Application feature.

About this task

Take these steps to copy or edit a launch definition:

Procedure

1. Select the Navigator item where the launch definition originates.
2. If the launch definition was created from a table, chart or event console view, open the workspace containing the view. Use the Workspace Gallery to see thumbnail graphics of the workspaces for a Navigator item, and select one to open.
3. Right-click one of the following:
   - Navigator item.
   - Row in a table or event console view.
   - Data point in a chart (such as a pie slice or bar).
   - Topology view object
4. Click ▶ **Launch** and select the definition you want to copy or edit.
5. If you want to create a copy of the definition, click **Create Another**.
6. Starting at the step for selecting the application to launch in the **Target** text box, continue with the remaining steps for "Creating a launch definition".

Deleting a launch definition

When you no longer need to keep a launch definition, remove it from the list.

Before you begin

Your user ID must have View and Modify permission for the Launch Application feature.

About this task

Take these steps to delete a launch definition:

Procedure

1. Select the Navigator item where the launch definition originates.
2. If the launch definition was created from a table, chart or event console view, open the workspace containing the view.
3. Right-click one of the following:
   - Navigator item.
   - Row in a table or event console view.
   - Point in a chart (such as a pie slice or anywhere in a gauge)
   - Topology view object
4. Click ▶ **Launch**.
5. Select the definition from the list and click **Delete**.
6. Click **Apply** to save your changes and keep the window open, or click **OK** to save your changes and close the window.

Launch application examples

These examples illustrate some of the uses for launch application and give the definition syntax.

Launching a batch file

To launch a batch file, you must specify the command prompt as the target and the batch file as the argument. This example starts myscript.bat.

Target: `c:\windows\system32\cmd.exe`
Arguments: `/c start c:\scripts\myscript.bat`

Start in: `c:\windows\system32`

Launching a website with search text

In this example the user typed everything in the **Arguments** text box except for &NT_Logical_Disk;:Disk_Name, which was added when DiskName was selected from the list of Substitutable items. The attribute is replaced by the current value when the application is launched.

Target: `C:\Program Files\Mozilla Firefox\firefox.exe`
Arguments: `http://www.ibm.com/search?q=+`
`&NT_Logical_Disk.Disk_Name`

Launching a website with no target browser specified

This example specifies a URL for the argument, much like the previous example. Notice, however, that there is no specified target. When the Target field is left blank, the Microsoft Internet Explorer is assumed (unless another browser was specified with the kjr.browser.default client parameter).

Arguments: `http://www.ibm.com`

Launching another Tivoli Enterprise Portal browser client on a different portal server This example is similar to the previous one. You can start a new browser client in another Tivoli managed network. If the browser supports multiple tabbed windows, a new tab is started for the launched portal workspace.

Target: `C:\Program Files\Mozilla Firefox\firefox.exe`

Arguments: `http://floor3west.usca.ibm.com:15200/`
`cnp.html?hostname=floor3west*`

Launching a Telnet session

To launch a Telnet session, you must specify the command prompt as the target and the Telnet session as the argument.

Target: `c:\windows\system32\cmd.exe`
Arguments: `/c start telnet.exe computer, ,` where computer is the location where the Telnet executable is found.
Start in: `c:\windows\system32`

Chapter 12. Using Tivoli Enterprise Portal user authorization

Every portal work session begins with a successful logon and connection to the Tivoli Enterprise Portal. The logon user IDs and user groups are created and profiled through the Administer Users window.

Administer Users is a multi-tabbed two-paned window. The top frame has two tabs: 🔲 **Users** and 🔵 **User Groups**, that list the user IDs, distinguished names if the portal server is configured for authentication to an LDAP user registry, and the user groups that are stored on the portal server. The profile of the selected user or user group is reflected in the bottom frame:

🔑 **Permissions** has a list of the portal features in the Authorities box. On the right are the possible operations for the selected feature. A selected check box means the selected user or user group has permission to perform that operation; a ▫ indicator next to the check box means the permission was added to a user group the user belongs to.

🔲 **Applications** shows all the applications being monitored and that are available for assigning to the user or user group. One user or user group, for example, can be profiled to see only the OMEGAMON applications, another to see only Linux and Oracle, middleware, and another to see all applications.

🔀 **Navigator Views** shows all the Navigator views that are on the portal server and that are available for assigning to the user or user group. The user or user group can be restricted to seeing only a certain branch of a Navigator view rather than the entire hierarchy.

🔵 **Member of**, when the Users tab is selected, or 🔲 **Members**, when the User Groups tab is selected, is a list of the groups the user belongs to or the user names in the group.

The User Administration function enables you to maintain user IDs and user groups on the portal server, and provides varying degrees of access to the features and views of your monitored environment to accommodate any combination of job roles, such as *operators* who respond to alerts and direct them to the appropriate person for handling and *administrators* who plan, design, customize, and manage the monitoring environment.

In some managed enterprises one person might assume all of these roles. In larger enterprises, the roles are often divided. You can choose to assign roles by individual user or by user type or both.

Tivoli Enterprise Portal user IDs are also required for users who access monitoring dashboards in IBM Dashboard Application Services Hub. How you manage dashboard users depends on the type of authorization configured in the portal server and whether the dashboard users will also use the Tivoli Enterprise Portal client. There are two types of authorization that can be configured for controlling access to monitored resources in IBM Dashboard Application Services Hub:

Role-based authorization policies
These policies are created using the tivcmd Command-Line Interface for Authorization Policy. They provide more granular authorization than Tivoli Enterprise Portal monitoring application assignments. Using role-based authorization policies, you can assign a user permission to view specific managed system groups or managed systems. When role-based

authorization policies are enabled in the portal server, dashboard users need a Tivoli Enterprise Portal user ID but do not require any Tivoli Enterprise Portal permissions or monitoring application assignments unless they are also Tivoli Enterprise Portal client users. In this case, role-based authorization policies control what resources they can access in the monitoring dashboards, and Tivoli Enterprise Portal permissions and monitoring application assignments control what they can access in the Tivoli Enterprise Portal client.

Tivoli Enterprise Portal authorization

This is the default authorization mechanism for dashboard users. A dashboard user must have a Tivoli Enterprise Portal user ID and be assigned the permissions and monitoring applications to control their access to resources in monitoring dashboards. If a dashboard user is also a Tivoli Enterprise Portal client user then they are assigned a single set of permissions that control what monitored resources they can access in both applications.

Configuring the portal server and Dashboard Application Services Hub to share an LDAP user registry is the best practice approach for having a federated set of dashboard users and Tivoli Enterprise Portal client users. In this scenario, the dashboard users login to the dashboard hub with their LDAP username and you must map their LDAP distinguished name to a Tivoli Enterprise Portal user ID with the required permissions.

Tivoli Enterprise Portal user IDs are automatically created with no permissions if a dashboard user requests monitoring data and does not have a user ID mapped to their distinguished name. See "Notes on user administration" on page 379 for more details.

Managing user IDs

Managing user IDs begins with planning the authorities to grant to users and whether they will belong to user groups.

The Administer Users window provides the tools for creating and maintaining user IDs, and adjusting permissions. This is also where user IDs are mapped to their unique identifier in the LDAP user registry if user authentication through the portal server has been configured.

Viewing and editing a user ID

After a user has been added to the **Users** list in the Administer Users window, you can check and edit the profile settings at any time.

Before you begin

To use this function, your user ID must have Modify permission for User Administration.

About this task

Use the following steps to edit a user ID:

Procedure

1. Click 🔸 **Administer Users**.
2. Do one of the following in the **Users** list:
   - Click inside the **Name** or **Description** field to edit either of them.
   - Double-click anywhere in a row to open the Modify User window for editing any of the fields.
   - Right-click the user profile you want to edit and click 🔸 **Modify User**.
3. Edit the **User Name**, **Distinguished Name** or **User Description**, then click **OK**. Distinguished Name is required if user authentication is through the portal server to an LDAP user registry. You cannot change the one-word User ID other than to change the letter casing. To edit the one-word User ID, delete the user profile and create a new one.
   - If you have not yet added the DN, click **Find** to locate the name that matches the user ID.
     If your monitored environment was previously configured for authentication through the Tivoli Enterprise Monitoring Server and then reconfigured to authenticate through the Tivoli Enterprise Portal Server, you might see two entries for the name. Select the one where o=defaultWIMFileBasedRealm and not O=DEFAULTWIMITMBASEDREALM.
4. To change the 🔑 Permissions, select a function from the **Authorities** tree and select or clear each option as appropriate for all functions with permissions that you want to change. You can change your own user permissions except Create and Modify for User Administration
5. To assign access privileges to applications (managed system types), click the ☐ Applications tab, select any applications you want to remove from the **Allowed Applications** list and ▶ move them to the **Available Applications** list; select the applications you want to add from the **Available Applications** list (or select **<All Applications>**), and ◀ move them to the **Allowed Applications** list.

   After selecting the first application, you can use Ctrl+click to select other applications or Shift+click to select all applications between the first selection and this one.
6. To change any Navigator view assignments, click the 🔗 Navigator Views tab, then add or remove Navigator views from the **Assigned Views** list, and select and ▲ move the one to be the default to the top of the list. For each Navigator view, change the **Assigned Root** as needed.
7. When you are finished editing the user profile, save your changes with **Apply** if you want to keep the Administer Users window open, or **OK** if you want to close it.

Results

The next time the user logs on, the permission changes will be in effect.

Adding a user ID

Create a user ID for all users that should be able to log onto the Tivoli Enterprise Portal Server using a portal client or the `tacmd tepsLogin` command. A user ID is also required for IBM Dashboard Application Services Hub users who request monitoring data. You can use the default user profile or copy the profile of an existing user.

Before you begin

To use this function, your user ID must have Modify permission for User Administration.

Procedure

1. Click 👤 **Administer Users**.
2. Create a new user ID or create one from another:
   - To create a new user ID with the default user profile, click 🗋 **Create New User**.
   - To create a new user ID from an existing one, select the profile that you want to use from the **Users** list and click 🗐 **Create Another User**.
3. In the Create New User window, enter the user information:
   - **User ID**: The logon name. The name must use ASCII characters, can be up to 10 characters, and can contain no spaces. The name is limited to eight characters if user authentication is at the hub monitoring server and uses RACF® (resource access control facility) security for z/OS.
   - **User Name**: The name of the user or job classification or both. This name can include spaces and be up to 32 characters. The user name is displayed in **Users** list.
   - **Distinguished Name:** The unique identifier in the Lightweight Directory Access Protocol (LDAP) user registry for the name given in the **User ID** field. Click **Find** to locate and insert the distinguished name, such as UID=FRIDA,O=DEFAULTWIMITMBASEDREALM
   - **User Description:** Optional description for the user. The text can include spaces and punctuation.
4. Click **OK** to close the window and see the new user ID arranged alphabetically in the **Users** list.
5. To change the 🔑 **Permissions**, select a function from the **Authorities** tree and select or clear each option as appropriate for all functions with permissions that you want to change.
6. To assign access privileges to applications (managed system types), click the ☐ **Applications** tab, then select **<All Applications>** or the individual applications the user should see, and click ◀ to move them to the **Allowed Applications** list. After selecting the first application, you can use Ctrl+click to select other applications or Shift+click to select all applications between the first selection and this one.
7. To assign Navigator views, click the 🔗 **Navigator Views** tab:
   a. Select a Navigator view (or more with Ctrl + click and Shift + click) from the **Available Views** and click ◀ to move it to the **Assigned Views**.
   b. Use ▲ to place the view that you want to be the default at the top of the list; use ▼ and ▲ to arrange the other Navigator views in the order that they should appear in the Navigator toolbar View ▾ list.
   c. For the selected Navigator view, change the **Assigned Root** as needed.
8. When you are finished creating the user profile, save your changes with **Apply** if you want to keep the Administer Users window open, or **OK** if you want to close it.

What to do next

The Tivoli Enterprise Portal client logon window has a field for entering a user ID and password. If you want the user ID and password to be authenticated,

configure the monitoring server or portal server to authenticate users. See Enabling user authentication for details.

Removing a user ID

You can remove a user ID as needed.

About this task

👤 To use this function, your user ID must have Modify permission for User Administration.

Use the following steps to remove a user ID:

Procedure

1. Click 👤 **Administer Users**.
2. Select the user ID that you want to delete. You can select additional user IDs with Ctrl+click, or with Shift+click to select all user IDs between the first selection and this one.
3. Click ✖ **Remove Users** to delete the selected user ID and profile from the list.
4. When a message asks you to confirm the user ID removal, click **Yes**. The user is permanently removed from the user ID list. If the user is currently signed on, this does not affect their work session, but they will not be able to log on again.

   Note: You cannot remove your user ID (the one you logged on with) or the <Default User> ID.

Default user

The first user ID in the **Users** list is <Default User>.

👤 To use this function, your user ID must have Modify permission for User Administration.

The Default User ID is used as the template ID for users created with 🗋 **Create New User**. Edit this user ID if you want to change any of the default settings. The initial defaults enable all the functions listed under Tivoli Enterprise Portal Authorities except User Administration Create and Modify. Any changes you make to the <Default User> ID apply to users created from this point on; they do not affect any existing user ID settings.

Managing user groups

User groups enable the administrator to authorize the same set of functional permissions, applications, and Navigator views to multiple users at one time. Management of user authorization can be done by groups as well as individually.

A user can be associated with one or more user groups. If a permission is granted to a user directly through their user ID, they maintain that permission even if a user group they belong to does not grant that permission. The reverse is also true, so that if an individual user ID is not granted a permission but the group ID is, the user will have the permission through their membership in the user group. Thus, the user's permission set is collected from what is given to the individual user ID and to any and all user groups that they belong to.

Authorization will also be by global authority and by association with managed system and managed system groups. This security is not dependent on external authorization.

When the active top tab is 👤 **Users**, the last tab on the bottom set of tabs reads 👥 **Member Of**. When the active top tab is 👥 **User Groups**, you will also have a 👤 **Members** tab. Assignment of users to groups can be done in either of these lower tabs.

Click the group in the details view at the top, then go to the 👤 **Members** tab to see the list of users that belong to this group. likewise, to see the groups a user belongs to.

Viewing user group memberships

You can view both the groups a user ID belongs to, and the list of user IDs belonging to a user group.

About this task

👤 To use this function, your user ID must have Modify permission for User Administration.

Procedure

1. Click 👤 **Administer Users**. The Administer Users window is divided into two, with **Users** and **User Groups** tabs at the top, and **Permissions**, **Applications**, **Navigator Views**, and **Member Of** below.
2. To see the groups a user belongs to, select a name from the 👤 **Users** list, then click the 👥 **Member Of** tab. The groups the user belongs to are listed in the **Assigned Members Of** list.
3. To see the user IDs assigned to a group, select a name from the 👥 **User Groups** list, then click the 👤 **Members** tab. The users belonging to the group are in the **Assigned Members** list.

Adding a user group

You can create a new user group from the beginning or you can copy a group with similar permissions and user assignments to what you want, then modify the copy.

Before you begin

To use this function, your user ID must have Modify permission for User Administration.

About this task

Complete these steps to add a user group:

Procedure

1. Click 👤 **Administer Users** to open the Administer Users window.
2. Click the 👥 **User Groups** tab.
3. Do one of the following:
   - To create a new user group, click 🗋 **Create New Group**.
   - To copy an existing user group, select the group name from the list and click 🗋 **Create Another Group**.

4. In the Create New Group or Create Another Group window, enter the following user information:

   a. **Group ID:** The group identifier. This name can be up to 10 characters and can contain no spaces. The name is limited to eight characters if the hub monitoring server uses RACF (resource access control facility) security for z/OS.

   b. **Group Name:** The name or job classification for the user group. This name can include spaces..

   c. **Group Description:** The text to describe the user group, such as their responsibilities. The description can include spaces and punctuation.

5. Click **OK** to close the window and see the new user group arranged alphabetically in the User Group list.

6. Add members to the group in the Members tab by selecting one or more user IDs in the **Available Members** list and clicking ◄ to move to the **Assigned Members** list.

7. To change the Permissions for the group, select a function from the **Authorities** tree and select or clear each option check box for all functions.

8. To assign access privileges to applications (managed system types) for the group, click the Applications tab, then select **<All Applications>** or the individual applications the user should see, and click ◄ to move them to the **Allowed Applications** list. After selecting the first application, you can use Ctrl+click to select other applications or Shift+click to select all applications between the first selection and this one.

9. To assign Navigator views to the group, click the Navigator Views tab, then add or remove Navigator views from the **Assigned Views** list, and use ▲ to place the default view at the top of the list. For each Navigator view, change the **Assigned Root** as needed.

10. When you are finished creating the user group, save your changes with **Apply** to keep the Administer Users window open, or **OK** to close it.

Reviewing and editing a user group

After a user group has been added to the **User Groups** list in the Administer Users window, you can check and edit the profile settings at any time.

About this task

To use this function, your user ID must have Modify permission for User Administration.

Use the following steps to edit a user ID:

Procedure

1. Click Administer Users to open the Administer Users window.

2. Click the User Groups tab.

3. Right-click the user group to edit and click .

4. Edit the **Group Name** and **Group Description**, then click **OK**. You cannot change the one-word group ID. You must, instead, create another user group from this one and give it a new name, then delete this one.

5. To change the Permissions, select a function from the **Authorities** tree and select or clear each option as appropriate for all functions with permissions that should change.

6. To change the group access privileges to applications (managed system types), click the ☐ Applications tab, select any applications you want to remove from the **Allowed Applications** list and click ➡ ; select the applications you want to add from the **Available Applications** list (or select **<All Applications>**), and click ◀ . After selecting the first application, you can use Ctrl+click to select other applications or Shift+click to select all applications between the first selection and this one.

7. To change any Navigator view assignments for the group, click the ☑ Navigator Views tab, then add or remove Navigator views from the **Assigned Views** list, and use ▲ to place the one you want to be the default at the top of the list. For each Navigator view, change the **Assigned Root** as needed.

8. When you are finished editing the user group, save your changes with **Apply** to keep the Administer Users window open, or **OK** to close it. The user group changes are effective the next time each group member logs on.

   Note: You can change the permissions, except Create and Modify for User Administration, of any groups you are a member of.

Removing a user group

You can remove a user group.

About this task

👤 To use this function, your user ID must have Modify permission for User Administration.

Use the following steps to remove a user ID:

Procedure

1. Click 👤 **Administer Users** to open the Administer Users window.
2. Click the 🌸 **User Groups** tab.
3. Select the user group to delete from the list and click ☑ **Remove Selected Group**. You can select additional user IDs with Ctrl+click, or with Shift+click to select all user groups between the first selection and this one.
4. When a message asks you to confirm the user group removal, click **Yes**. The group is permanently removed from the user group list. Any members of this user group who receive permissions from the group will not be affected until they next log on to the portal server.

Administer Users

Your user ID and the user groups you are a member of are profiled with a set of permissions that determines which Tivoli Enterprise Portal features you are authorized to see and use, a list of monitored applications you are authorized to see, and a list of Navigator views (and the highest level within a view) you can access.

Clicking 👤 **Administer Users** opens the Administer Users window. This is a two-paned window with Users and User Groups tabs in the top frame, and several tabs in the bottom frame. This arrangement enables the administrator to manage user profiles by individual user, by user groups, or a combination of the two. You might create a user profile, then copy the profile for each additional user and change settings as needed (such as, for the Action feature, granting View permission to one user and granting Modify permission to a different user). Or

you might create a user group with a particular profile and add users to the group. Then you can modify the permissions once for a group and apply to all members automatically.

When you modify the permissions or the list of monitored applications for a Tivoli Enterprise Portal user or user group, the authorization change does not take effect until the user logs out of the Tivoli Enterprise Portal client and then logs back in. If Tivoli Enterprise Portal permissions are being used to authorize monitored resources for dashboard users, the authorization changes also do not take effect until the user logs out of Dashboard Application Services Hub and logs back in.

Users and User Groups

The ⚇ **Users** and ⚇ **User Groups** tabs list the user IDs and the user groups that are stored on the portal server.

After you select a user or user group from one of the lists, you can click any of the tabs in the lower half of the window to see the what permissions have been granted and what has been assigned. User groups enable the administrator to authorize the same set of functional permissions, applications, and Navigator views to multiple users at one time. Management of user authorization can be done by groups as well as individually. A user can be associated with one or more user groups; authorization by group will be by inclusion and not exclusion (nested groups are supported). Authorization will also be by global authority and by association with managed system and managed system groups. This security is not dependent on external authorization.

Permissions

You can authorize the same set of functional permissions multiple users, user group or each user ID at one time.

The following features are enabled or disabled individually for each user ID or user group.

Action

> ☑ **View** allows the user to see and run a take action command from the available list of commands in the Take Action view and in the pop-up menu for the Navigator item.

> ☑ **Modify** allows the user to create and save Take Action commands. When enabled, ⚡ **Edit Action** appears in the Navigator pop-up menu.

> When issuing a take action command, you must be authorized on the relevant system for the requested command. For example, to issue a TSO command, your user ID must be both a valid TSO ID and a valid user ID on the portal server. The user ID must be typed with the same letter casing exactly as typed when logging on to the portal server (with the same letter casing).

Agent Management

> ☑ **Manage** allows the user to perform agent deployment throughout the managed network. This includes installing a monitored product, keeping the software revisions up-to-date, and removing an agent from the managed network. This permission also requires Action - Modify to be enabled.

> ☑ **Start/Stop** allows the user to start a monitoring agent or to stop it running.

Custom Navigator Views

☑ **Modify** allows the user to create new Navigator views, edit and delete them. With Modify cleared, the user will not see ☑ **Edit Navigator View** in the Navigator toolbar.

Event ☑ **Attach** allows the user to attach a file (such as detailed notes) to the situation event. This permission requires that the user also have the Acknowledge and View permissions.

☑ **Close** lets you close a pure event or an event that was open before a situation was stopped manually. When it is enabled, ☑ **Close Situation Event** appears in the pop-up menu of the situation event flyover list, event Navigator item, and situation event console view when the selected event is a pure event or the situation has been stopped.

☑ **View** enables you to see situation event indicators in the Navigator when situations become true.

☑ **Acknowledge** allows you to acknowledge a situation event. When this permission is enabled, **Acknowledge Event** appears in the pop-up menu of the situation event flyover list, event Navigator item, and situation event console view.

Feature

☑ **Enable** is dimmed because you cannot change it. The access to this feature is determined by your organization's IBM Tivoli Monitoring license.

History

☑ **Configure** allows the user to open the History Collection Configuration window, configure history files and data rolloff, and start and stop data collection for the different attribute groups. When this permission is enabled, ☑ **History Configuration** appears in the main toolbar.

Launch Application

☑ **Launch** allows the user to invoke any of the launch definitions available for the Navigator item, table view, chart view, or situation event console view. When this permission is enabled, ☑ **History Configuration** appears in the main toolbar.

☑ **View** allows the user to see the composition of the selected launch definition.

☑ **Modify** allows the user to create, edit and delete launch definitions.

Managed System Group

☑ **View** allows the user to access the Object group editor for viewing managed system groups. The user also needs Modify permission for the Object group editor tools to be available.

☑ **Modify** allows the user to open the Object group editor to create, edit, and delete managed system groups.

Policy ☑ **View** allows the user to open the Workflows window to see the policies and their definitions. With View permission, the ☑ **Workflow Editor** is available in the main toolbar and ☑ **Manage Policies** is available in the Navigator pop-up menu at the ☑ agent level.

☑ **Start/Stop** lets you start and stop policies. With this permission enabled, ▶ **Start Policy** and ● **Stop Policy** are available when you select a policy.

☑ **Modify** allows the user to open the Workflow editor to create and edit policies. With the Modify permission enabled, ▯ **New Policy** is available after the user selects a policy, as are the other editing tools: ▧ **Edit Workflow**, ▯ **Copy Policy**, and ▨ **Delete Policy**.

Query ☑ **View** allows the user to access the Query editor through the Properties editor and select a query for the selected table or chart. With the View permission enabled, the user can assign a query through the Query tab of the Properties editor.

☑ **Modify** allows the user to create, edit and delete queries in the Query editor. With the Modify permission enabled, ▯ **Query Editor** is available from the main toolbar, as are the query editing tools.

Situation

☑ **View** allows the user to see situations in the Situation editor, including any expression overrides, and in the Manage Situations at Managed System window. With the View permission enabled, ▯ **Situation Editor** is available in the main toolbar and in the Navigator item (except at the platform level) pop-up menu.

☑ **Modify** lets you create new situations and manage them. When the Modify permission has been granted, the situation editing tools and pop-up menu options are available in the Situation editor, as well as the **Override Formula** button in the Distribution tab for qualifying situations.

☑ **Start/Stop** lets you start or stop a situation and enable or disable a situation override. When this permission is enabled, ▶ **Start Situation** and ⏺ **Stop Situation** are available in the situation event flyover list, situation event console view, Situation editor pop-up menu, and the Manage Situations at Managed System window; and ⊘ **Enable Situation Overrides** and ⊖ **Disable Situation Overrides** are available in the Situation editor pop-up menu.

Terminal Script

☑ **View** allows the user to run or stop running a terminal emulator script and to see scripts, but not to edit them. If View is disabled the user will be able only to run or stop a script.

☑ **Modify** allows the user to create or record new terminal emulator scripts, edit, and delete them.

User Administration

If you are viewing your own user ID, you will see that View and Modify are disabled; you cannot change your User Administration permissions.

☑ **Logon Permitted** enables log on to the portal server with this user ID. The administrator can clear the check box to deny a user access to the portal. This option works in conjunction with the KFW_AUTHORIZATION_MAX_INVALID_LOGIN (the default is 0, unlimited attempts are allowed) parameter in the Tivoli Enterprise Portal Server Environment Configuration file, *kfwenv*. When the value has been set and the invalid attempts have been exceeded, the check box is cleared automatically and the administrator must select the check box to reset the logon attempt count. See the *IBM Tivoli Monitoring Administrator's Guide* for details. ☐ **Modify** allows the editing of user IDs and removing them.

When this permission is enabled, ▯ **Administer Users** is available in the main toolbar and the tools are available in the Administer Users window.

☑ **Author Mode Eligible** allows the user to enable or disable their Author Mode permission under **Workspace Administration** (see next authority), but not for any other user IDs.

☐ **View** allows the user to open the Administer Users window and see their user profile.

☑ **Administration Mode Eligible** allows the user to enable or disable their Administration Mode permission under **Workspace Administration** (see next authority), but not for any other user IDs.

Workspace Administration
☑ **Workspace Author Mode** allows the user to create and edit workspaces, links, and terminal emulator scripts. If Workspace Author Mode is disabled, the user cannot make any of these changes but can continue monitoring and responding to alerts; the tools can still be seen, but they are disabled.

☐ **Workspace Administration Mode** is available only for the SYSADMIN user ID and new IDs made from it in the Create Another User window. When administration mode is enabled, changes you make to workspaces affect all users who log on to the same portal server. When it is disabled, workspace changes you make are not shared with other users. Be sure to select ☑ **Do not allow modifications** in the Workspace Properties whenever you create or edit a workspace in administration mode. Otherwise, if a user edits that workspace, you no longer own the workspace and cannot override their changes.

WebSphere MQ Configuration Authorities
IBM Tivoli OMEGAMON XE for Messaging: WebSphere MQ Configuration installations will see this folder.

☑ **View** allows the user to see, but not change, your organization's WebSphere MQ configuration in the Navigator Configuration view.

☑ **Modify** allows the user to change your organization's WebSphere MQ configuration or to schedule updates in the Configuration view.

Storage Subsystem Authorities
IBM Tivoli OMEGAMON XE for Storage installations will see this folder. Select ☑ **View** to allow the user to see, but not change data. Select ☑ **Modify** to allow the user to change data.

Data Collection Configuration allows the user to view or modify the collection control interval for the DFSMSrmm Status workspace and Dataset Attributes System Summary workspace.

Dataset Groups Collection Interval allows the user to view or modify the control interval for the Dataset Group Summary workspace.

Define/Update Dataset Groups allows the user to view or modify the group definitions for the Dataset Group Summary workspace.

Applications

Your user ID is set so you can see some or all the application types being monitored. For example, one user might be able to see only mainframe applications, while another can see only middleware, and another sees all applications.

Allowed Applications
Shows the applications that you can access from Tivoli Enterprise Portal.

Available Applications

Shows the applications available for assignment to the selected user. If **<All Applications>** is on the **Allowed Applications** list, then no other entries can be added. You must move it back to **Available Applications** before you can select a subset of applications to assign.

Select the applications you want to add, or select **<All Applications>**, and ◀ move them to the **Allowed Applications** list. After selecting the first application, you can use Ctrl+click to select other applications or Shift+click to select all applications between the first selection and this one.

Navigator views

When a Navigator view is created, only the author is able to see the view, but it is available for the administrator to assign to users. An assigned Navigator view means the user can open it. For each assigned view, the user can be restricted to see only a certain branch rather than the entire hierarchy.

Assigned Views

Shows the Navigator views the user is able to see and access. The first Navigator view in this list is the default for the user; it displays automatically whenever the user logs on. You can select any views to which you do not want the user to have access, and click ➡ right arrow to move them to the **Available Views** list. Select the appropriate entries and click ◀ left arrow to move them to the **Assigned Views.** You can move a Navigator view to the top of the list to make it the default by clicking the ▲ up arrow.

Available Views

Shows the Navigator views not assigned to the user and available for assignment. Select the Navigator views you want to add and move them to the **Assigned Views** list by using the ◀ left arrow. After selecting the first view, you can use Ctrl+click to select other views or Shift+click to select all views between the first selection and this one.

Assigned Root

Shows the Navigator view chosen in Assigned Views, with the user's assigned Navigator root highlighted. The root is the top-most level of this Navigator view that the user can access. The user can access this item and all items below it, but no items parallel to or above it in the Navigator.

For example, you can assign UNIX Systems as the assigned root. The user sees the UNIX Systems workspaces and those below, but is not able to see the Enterprise workspaces or anything under Windows Systems.

Member Of and Members

When you select a user or user group from the list, the last tab on the bottom set of tabs reads either **Member Of** or **Members** (reflecting the selection of a User or User Group). Assignment of users to groups can be done in either tab.

Notes on user administration

Read these notes to understand the user ID contribution to Tivoli Enterprise Portal functions and modes.

Workspace administration mode

Any changes you make to workspaces, links, and terminal host session scripts in the Tivoli Enterprise Portal are available only to your user ID. The exception is while Workspace Administration Mode is enabled.

Workspace administration mode enables you to customize and add workspaces, links, and terminal emulator scripts that are shared with all users connected to the same Tivoli Enterprise Portal. See "Starting workspace administration mode" on page 69.

SYSADMIN logon ID

The Tivoli Enterprise Portal requires your logon ID whenever you start a work session. Every ID must first have been registered on the portal server. You can log onto the portal server with **SYSADMIN** and register other user IDs through the Administer Users window. The initial user ID, **SYSADMIN**, has full access and complete administrator authority. The system administrator registers additional users and sets their access privileges and authority.

User ID and groups

Each user ID is stored at the Tivoli Enterprise Portal Server and contains:
- The user name
- Job description
- Permissions for viewing or modifying Tivoli Enterprise Portal functions
- Assigned Navigator views and which Navigator item in each view appears as the root (default is the first item)
- Access to specific monitoring applications
- The user groups the user belongs to and indicators to signify when a permission has been granted to the user by a user group

Each user group is also stored at the portal server and has the same contents as for individual user IDs. But, instead of a list of user groups, it has a list of the user IDs assigned to the group.

Default user

The first user ID in the list is **<Default User>** and is used as the template ID for users created with Create New User. Edit this user ID if you want to change any of the default settings. The initial defaults enable all the functions listed under Tivoli Enterprise Portal Authorities, except the Modify permission for **User Administration**. Any changes you make to <Default User> ID apply to users created from this point on; they will not affect any existing user ID settings.

Granting access to a user

You set the authority privileges for each user when you create their user IDs. Giving users access to operational areas and customization options takes planning. Consider the job responsibilities of each user and the company security requirements when specifying authority privileges.

Important: Anyone with permission to create custom queries obtains access to all of the ODBC data source names (DSNs) created at the Tivoli Enterprise Portal Server. Add database user IDs, to be used in the DSN, to your database software,

making sure to restrict user access to only those tables, columns, and so on, allowed by your organization's security policies.

Automatic Tivoli Enterprise Portal user ID creation

The first time a new user accesses a monitoring dashboard in IBM Dashboard Application Services Hub, a Tivoli Enterprise Portal user ID is automatically created and mapped to the user's LDAP distinguished name if a Tivoli Enterprise Portal user ID does not already exist for the user. The Tivoli Enterprise Portal user ID is a randomly generated string. If you need to assign Tivoli Enterprise Portal permissions and monitoring applications to a dashboard user and their Tivoli Enterprise Portal user ID was automatically created, you can either assign the permissions to the randomly generated user ID or perform these steps:

1. Delete the Tivoli Enterprise Portal user ID that was automatically created.
2. Create a new user Tivoli Enterprise Portal user ID, map it to the LDAP distinguished name for the user, and then assign it permissions and monitoring applications.

Validating user access

The Tivoli Enterprise Portal Server verifies user IDs whenever users log on. If a job description changes and the user requires different access to the portal server, you must review and perhaps change the user's permissions.

The user ID for logging on to the portal server might include a password. You do not establish passwords in the portal. Instead, you must define a matching user ID with password to the network domain user accounts or to the operating system where the Tivoli Enterprise Monitoring Server resides:

- User Accounts on the Windows system
- Password file on the UNIX system
- RACF or ACF/2 host security system on the z/OS system

As well, the monitoring server must be configured to Validate User. When users log on to the portal server, the hub monitoring server makes a request to the domain or the operating system to validate the user ID and password.

If the monitoring server has been installed on a distributed system, you can check if it has been configured to Validate User:

1. Start the Manage Tivoli Enterprise Monitoring Services program:

   Windows Start → Programs → IBM Tivoli Monitoring → Manage Tivoli Enterprise Monitoring Services.

   UNIX Change to the *install_dir*/bin directory and run the following command: `./itmcmd manage [-h install_dir]` where *install_dir* is the installation directory (default is opt/IBM/ITM).

2. Right-click the Tivoli Enterprise Monitoring Server row for TEMS1 (hub) and select **Reconfigure**.
3. In the Tivoli Enterprise Monitoring Server Configuration window, observe the setting of the ☑ **Security: Validate User** check box.

   When this option is selected, the password is required whenever a user logs on to the portal server; when it is cleared, the user name is required to log on but no password is required.

Note: Be aware that passwords must follow the security requirements for your organization. If this includes periodic password changes, you might get a **Logon password has expired** message while attempting to log on to the portal server. Should this happen, you must change your system password before you can log on. For example, on Windows this means changing the password through the Administrative Tools User Accounts.

Launching into the portal from other applications

In addition to any security requirements for launching into the Tivoli Enterprise Portal (such as single sign-on requirements), the Tivoli Enterprise Portal user ID that receives control after a launch from an external application must be pre-authorized to access the target managed system and workspaces. The user ID also must be authorized to issue any required take action commands.

User ID for Take Action commands

When the Tivoli Enterprise Portal sends a Take Action command to a managed system, the user ID might or might not be checked for authority to perform the action. In the simplest case, the command is sent to the managed system and executed using the user ID under which the agent is running. TheTivoli Enterprise Portal user ID is sent along with the action command in these contexts:

- On-demand: user ID currently logged on
- Situation action: user ID of the last person to update the situation
- Workflow action: user ID of the last person to update the policy

However, the ID is ignored by the managed system unless told otherwise by a command prefix. These are command handlers implemented in the IBM Tivoli Monitoring products to control whether the Tivoli Enterprise Portal user ID should be validated before passing the command to the agent for execution.

Command prefix

When a command prefix is present in the Take Action, the agent passes the command to the application handler rather than executing the command. The syntax of the prefix and take action command is **productcode:CNPuserID:command** and the agent routes it to the application for execution. The application is free to execute the command with whatever user ID is appropriate. In the case of OMEGAMON XE for WebSphere MQ on z/OS, the Tivoli Enterprise Portal user ID is used.

If the special prefix is missing, the agent executes the command with the user ID under which the agent is running.

Most monitoring products do not employ a command prefix. IBM Tivoli Monitoring for WebSphere MQ does and, in fact, prepends any on-demand Take Action commands with a hidden **MQ:CNPuserID:** prefix, although you cannot see it.

UNIX setuid command

In addition to the command prefix and security exit, UNIX offers another option: a setuid command, which causes the process to dynamically change its userid. Thus, the agent could be changed to set the ID to the value passed as a parameter, issue the command, then change the user ID back again after the command is issued.

Troubleshooting logon error messages

Logon prompts and progress messages are displayed in the Logon window status bar. If a user cannot log on, a message is displayed.

If a user cannot log on, one of the following messages is displayed:

Failed connection to Tivoli Enterprise Portal Server

1. On the system where the Tivoli Enterprise Portal Server is installed, click **Start** -> **Programs** -> **IBM Tivoli Monitoring** -> ▓ **Manage Tivoli Monitoring Services**.
2. Optional: Right-click the Tivoli Enterprise Portal Server entry and click **Change Startup**. In the window that opens, select ◉ **System Account** and ☑ **Allow Service to Interact with Desktop** and click **OK**.

    This opens a command line window when the Tivoli Enterprise Portal Server is started and displays the internal commands.
3. Ensure that the Tivoli Enterprise Portal Server is started:

    If it is started, recycle it.

    If it is stopped, start it.
4. If you are still unable to connect, review the following information. If it does not apply to your situation, contact IBM Software Support.

    If you are running in browser mode and going across networks to reach the Tivoli Enterprise Portal Server, it is possible the host name cannot be resolved by the network system. If this is the case, doing the following should resolve the problem:

1. On the system where the Tivoli Enterprise Portal Server is installed, click **Start** -> **Programs** -> **IBM Tivoli Monitoring** -> ▓ **Manage Tivoli Monitoring Services**.
2. Right-click the Tivoli Enterprise Portal Browser service and click **Reconfigure.**
3. Change the host name to the IP address in two places:

    In the **Launch URL** field, change *hostname* in http://*hostname*:15200 to the IP address of the Tivoli Enterprise Portal Server. For example, `http://10.21.2.166:15200`.

    In the **Host** field, change the host name to the IP address of the Tivoli Enterprise Portal Server.
4. Click **OK**.
5. Start Tivoli Enterprise Portal browser mode using the IP address instead of the host name.
6. If you are still unable to connect, contact IBM Software Support.

Logon password has expired

If the hub Tivoli Enterprise Monitoring Server is set to Validate Users, then passwords are required. Passwords must follow the security requirements of your organization. If this includes periodic password changes, you might get this message while attempting to log on to the portal server. Should this happen, you must change your system password before you can log on. For example, on Windows this means changing the password through the Administrative Tools User Accounts.

User authorization has failed -OR- Unable to process logon request

Tivoli Enterprise Portal uses the TEPS database to locally validate users. If

your hub monitoring server is set for user validation (Windows default), the user ID is also validated at the monitoring server to verify the password.

The portal server did not validate the user credentials as entered. For the "Unable to process logon request" message, the portal server was able to validate the user credentials but did not complete the logon request. In either case, have the user try logging on again. If the message is displayed again, do the following:

1. On the system where the monitoring server is installed, ensure that the server is running in ▮ Manage Tivoli Monitoring Services.

2. If the monitoring server is running, ensure that the user ID has been defined in Tivoli Enterprise Portal: Click ▮ **Administer Users**, then find the ID in the **Users** list.

3. If the user has been defined, check if host level security was turned on for the hub monitoring server and that the user ID has been authorized to the host environment:

   In ▮ Manage Tivoli Monitoring Services, right-click **Tivoli Enterprise Monitoring Server**, and click **Reconfigure.** If host level security has been configured, the **Security: Validate User** box is selected.

   If the monitoring server has been configured to Validate User, the user ID for Tivoli Enterprise Portal must also be added to the network domain user accounts or to the operating system where the monitoring server is installed, including a password.

   If non-ASCII characters were included in the user ID, they are not saved with the user ID.

4. Try logging on to Tivoli Enterprise Portal with the user ID in question.

5. If you cannot log on to Tivoli Enterprise Portal and the monitoring server is running properly, the problem might be with the Tivoli Enterprise Portal Server. Try recycling the portal server. If the user is still unable to log on, contact IBM Software Support.

This message is also displayed after a retry period of several minutes (the default is 10 minutes and can be changed through ▮ Manage Tivoli Monitoring Services) where the status bar shows **Validating user credentials** continuously. This can be a symptom that the monitoring server is stopped.

Chapter 13. Managing the environment

Managing your Tivoli monitored environment involves monitoring and managing the availability of the components of the Tivoli Management Services infrastructure and the agents that use its services. It also involves organizing managed systems and situations into logical groupings that can be applied to situations and policy distributions with a single action.

These topics describe the tools and functions you can use to monitor the status of managed systems and manage components and agents on distributed systems. See *Administering IBM Tivoli Monitoring Products* for detailed information on day-to-day administrative tasks.

Grouping objects

Use the Object group editor to organize situations, managed systems, and historical configurations into named collections that can be applied with a single action.

Situations in groups can be started or stopped, and they can be distributed to managed systems. Managed systems can be assigned as a group to custom Navigator view items, to situation distribution lists, to policy distribution lists, and to historical collection configurations. Similarly, historical collections in groups can be distributed to managed systems as if they were a single collection.

Object group editor

Use the Object group editor to organize Tivoli Monitoring objects (historical configurations, managed systems, and situations) into named groups that you can reference in distribution lists.

The navigation tree shows the types of object groups that are available and any groups that have been created. Click ⊞ to expand the list of groups for the object type. Some object types have categories that expand to show their groups. To replace this user assistance with an object group definition, do either of the following:

- Click a ⊞ group name.
- Click a group type, click 🗋 **Create new group**, and enter a name and description.

The object definition has tabs that show the objects that are assigned to the group and those that are available. When a member of a group is an object group itself, you can select it to see its constituents. You can assign individual objects to a group and you can assign a group to another group, thus creating nested groups.

Historical configurations
> Use historical configuration groups to combine related historical collections into groups. You can then control collection for the group rather than having to select collection definitions individually.

> 🖿 **Members**
>> The text boxes show the historical configuration members of a group, as well as any groups that are contained in (members of)

other object groups, and any that are available for inclusion in a group. Individual members represent historical collection definitions for an attribute group.

⊞ **Member of**

This tab shows any groups that the selected historical configuration group belongs to and any that are available.

◉ **Distributions**

The distributions are the managed systems and managed system groups that the historical configuration group is distributed to, as well as any available managed systems and groups. Select a managed system group to see its constituent objects in the Selected Assigned Object Members or Selected Available Object Members list.

Managed system

Any set of managed systems that have a relationship, such as by geography or business function, can be organized into a managed system group for assigning to custom Navigator views, historical collection distributions, situation distributions, or policy distributions. The navigation tree lists all the managed system types that are available, including ◉ **All Managed Systems** for creating managed system groups that combine multiple managed system types.

◉ **Members**

The text boxes show the managed system members and groups that have been assigned to the group and any members that are available for assignment. Any members that are object groups themselves can be selected from the *Members...* list to see the objects that belong to it in the *Selected...* list below.

⊞ **Assigned to**

Any groups that the managed system group is assigned to are shown here. If, for example, a situation or situation group has been distributed to this managed system group, that distribution will be listed here.

Situation

If you frequently distribute a particular set of situations, create a situation object group for them. You can then distribute the group in the same way as for a single situation, thus saving time and ensuring consistency in applying your best practices.

▣ **Members**

The text boxes show the situation members and groups that have been assigned to the group and any that are available for assignment. If a member is an object group, you can select it from the Members list. The objects that belong to it are displayed in the corresponding Selected list below.

⊞ **Member of**

Shown here are any groups the situation group belongs to and any that are available.

◉ **Distributions**

The distributions are the managed systems and managed system groups that the situation group is distributed to, as well as the available managed systems and groups. You can select a managed

system group to see its constituent objects in the Selected Assigned Object Members or Selected Available Object Members list.

Managed system groups

A managed system group is a named collection of managed systems. When distributing a situation, policy, or historical collection configuration to managed systems, or when assigning managed systems to a custom Navigator view, you can specify a managed system group (or groups).

If the main toolbar has no ⊞ **Object group editor**, your user ID does not have View permission for managed system groups. If you can open the editor, but the tools not enabled after you select one of the managed system groups, your user ID does not have Modify permission for managed system groups.

Predefined managed system groups

The Tivoli Enterprise Monitoring Server and every IBM Tivoli Monitoring product has at least one predefined managed system group, indicated by an asterisk at the beginning of the list name, such as *NT_SYSTEM for IBM Tivoli Monitoring for Windows.

You can also create new managed system groups for different purposes. For example, you can create a group of Linux managed systems for a particular geographic region named LINUX_LONDON, so that you can distribute situations that are specific to that location. Managed system groups are stored at the Hub Tivoli Enterprise Monitoring Server, thus any lists you create are available to all users who are logged on to the Tivoli Enterprise Portal Server that connects to that hub monitoring server.

When you assign a predefined managed system group to a situation, policy, query, custom Navigator, historical collection, historical configuration object group, or a situation object group, all managed systems of that type are selected. And as new managed systems are added to the monitored network, they automatically become members of the predefined managed system group of their agent type.

You cannot edit or delete a predefined managed system group.

The monitoring server provides three managed system groups. These lists can be applied at the Enterprise level of the Navigator Physical view and anywhere in a custom Navigator view:

*ALL_CMS applies to the hub monitoring server and any remote monitoring servers.

*EIB applies to the hub monitoring server.

*HUB applies to the hub monitoring server.

Creating a managed system group

Organize managed systems into named groups that can be assigned to the distribution of situations, policies, historical collections, queries, and dynamic Navigator items in custom Navigator views.

About this task

Take these steps to create a managed system group:

Procedure

1. Click ⊞ **Object group editor**.

2. If ⊚ Managed system is collapsed, ⊞ expand it.

3. Click one of the available managed system types, then click 🗍 **Create new group**. One of the managed system types is ⊚ **All Managed Systems**. Select it if you want to combine multiple agent types (such as Windows OS and Oracle) in one managed system group.

4. Type a descriptive name for the managed system group and click **OK**. The new managed system group is displayed in the managed system folder.

5. Optionally, select a managed system from the **Available Managed Systems** list and ◀ move it to the **Assigned** list. You can select multiple managed systems by holding down Ctrl while clicking each managed system. You can also, after selecting a managed system, use Shift+click to select all managed systems between this and the first selection.

6. Save your changes and either keep the editor open with **Apply** or exit with **OK**. The managed system group is now available to all users who are signed on to the portal server that connects to this hub monitoring server.

   Note: Be sure that the managed systems added to the managed system group are at a supported version level. Otherwise, it is possible that after you distribute a situation to the managed system group and then attempt to start the situation, its status will show as 🔴 Problem in the situation event console and message log views.

What to do next

You can now assign the managed system group to any object that has managed system assignments. Each of these locations has a **Edit Managed System Groups** button for convenient access to the Object group editor for managed system groups:

- **Distribution** tab of the History Collection Configuration editor
- **Distribution** tab of the Situation editor
- **Monitored resources** tab of the Navigator item properties
- **Change Policy Distribution** dialog box, which opens after you select the ☑ **Distributed** check box for a policy in the Workflows editor
- **Query Results Source** tab of the Query Editor

Copying a managed system group

You can create a copy of a managed system group as needed.

About this task

Use the following steps to create a copy of a managed system group:

Procedure

1. Click ⊞ **Object group editor**.

2. Click ⊞ to expand a managed system folder.

3. Select the managed system group you want to copy.

4. Click **Create another group**.

5. Type a descriptive name for the managed system group and click **OK**. The new managed system group is displayed in the managed system folder.

6. Remove a managed system from the **Assigned** list by selecting it and clicking
   ▶ . Add a managed system to the list by selecting it in the **Available Managed
   Systems** list and clicking ◀ . You can select multiple managed systems by
   holding down Ctrl while clicking each managed system. You can also, after
   selecting a managed system, use Shift+click to select all managed systems
   between this and the first selection.

7. Save your changes and either keep the editor open with **Apply** or exit with
   OK. The managed system group is now available to users whose Tivoli
   Enterprise Portal Server connects to the same hub monitoring server.

Editing a managed system group

You can edit a managed system group to better suit your environment.

About this task

Use the following steps to edit a managed system group. Be aware that any
situations, policies, queries, or custom Navigator views this managed system group
is assigned to might be affected by changes you make to the list.

Procedure

1. Click ⊞ **Object group editor**.
2. Click ⊞ to expand a managed system folder.
3. Select the managed system group that you want to edit.
4. To remove a managed system from the **Assigned** list, select it and ▶ move it
   back to the Available list. To assign a managed system, select it in the **Available
   Managed Systems** list and ◀ move it to the Assigned list. You can select
   multiple managed systems by holding down Ctrl while clicking each managed
   system. You can also, after selecting a managed system, use Shift+click to select
   all managed systems between this and the first selection.
5. Save your changes and either keep the editor open with **Apply** or exit with
   OK.

Deleting a managed system group

You can delete a managed system group if it is no longer needed.

About this task

Be aware that if you delete a managed system group that is assigned to a situation,
policy, query, or custom Navigator view, the assignment will no longer be valid.

To delete a managed system group:

Procedure

1. Click ⊞ **Object group editor**.
2. Click ⊞ to expand a managed system folder.
3. Select the managed system group you want to delete and click **Remove
   selected group**. You can also right-click the managed system group and click
   Remove selected group.
4. When a message asks if you are sure, click **Yes**; or click **No** if you are not sure
   of the effect this will have on your monitoring configuration.
5. Select another managed system type or list to work with or click **Cancel** to
   close the editor.

Creating a historical configuration group

Organize the historical collections that you have defined into named groups that can be managed collectively.

Before you begin

Historical configuration groups are made up of historical data collections that were defined in the History Collection Configuration window. Create the historical collections that you intend to include in the historical group before opening the Object Group editor.

This feature is available when the Tivoli Enterprise Portal Server and Tivoli Enterprise Monitoring Server are at Version 6.2.2 or later and the historical data collection that you want to add to the group is configured for **Managed System (Agent)** as the distribution method.

About this task

Take these steps to create a historical configuration group:

Procedure

1. Open the Object group editor from either of these locations:
   - Click ⊞ **Object group editor**, then click ⚏ **Historical configurations**.
   - If the History Collection Configuration window, is not open, click ⚏ **History Configuration**. Click a ⊚ monitored application, click an ▬ undistributed or ⚏ distributed historical collection definition, then click **Group**.
2. Click ⬚ **Create a new object group**.
3. Type a name for the historical group and, optionally, a description, then click **OK**. The new historical configuration group is displayed in the ⚏ **Historical configurations** branch.
4. In the **Members** tab, select a historical collection from the list of **Available Members** and ◂ move it to the **Assigned Members** list. You can select multiple collection definitions for an attribute group by holding down Ctrl while clicking; or use Shift+click to select all entries between this and the first selection. Historical collections that are configured for "Managing System (TEMS)" distribution are not displayed in the list.
5. If you have historical configuration groups available and want to assign a group to this one, click the **Member of** tab, then select a group and ◂ move it to the **Assigned Members Of** list.
6. In the **Distributions** tab, select the managed systems and managed system groups to assign to the historical configuration group. Select a managed system group to see its constituent objects in the **Selected Assigned Object Members** or **Selected Available Object Members** list.
7. Save your changes and either keep the editor open with **Apply** or exit with **OK**.

Results

The historical configuration group is now available to all users who log onto the same portal server. Historical data samples are collected to the short-term history files for the managed systems in the distribution.

When you next view the historical data collection in the History Collection Configuration window, the Distribution tab shows the historical configuration group in the **Start collection on** list. The group name is disabled (dimmed); you cannot remove it from the distribution. To do that, you would return to the Object group editor and remove the historical collection from the historical group.

Related tasks:

Creating a managed system group
Organize managed systems into named groups that can be assigned to the distribution of situations, policies, historical collections, queries, and dynamic Navigator items in custom Navigator views.

Creating a situation group

Organize situations that you have defined into named groups that can be managed collectively. The situation group can include a distribution of managed systems.

Before you begin

Your use ID must have Modify permission for Situations to use this function.

About this task

You can start from the Situation editor or from the Object group editor. Take these steps to create a situation group:

Procedure

1. Create a situation group from either the Situation editor or Object Group editor:
   - Start from the Situation editor:
     a. Open or create a situation that you want to add to the new situation group.
     b. While the situation is open in the Situation editor, click **Group**, then ◉ **Create a new object group**.
     c. Enter a name for the new situation group and, optionally, a description. The name can be up to 128 characters.

     After you click **OK**, the Object group editor opens the new situation group with the situation listed in **Assigned Members**.
   - Start from the Object group editor:
     a. Click ⊞ **Object group editor**.
     b. Click ▣ **Situation**, then ▯ **Create a new object group**.
     c. Enter a name for the new situation group and, optionally, a description. The name can be up to 128 characters.

     The new situation group is displayed in the ▣ **Situation** branch of the tree.
2. In the ▣ **Members** tab, select a situation from **Available Members** and ◀ move it to the **Assigned Members** list. You can select multiple situations by holding down Ctrl while clicking; or use Shift+click to select all entries between this and the first selection.
3. If you have situation groups available and want to assign a group to this one, select a group and ◀ move it to the **Assigned Members** list. You can also assign situation groups from the ⊞ **Member of** tab.
4. In the **Distributions** tab, select the managed systems and managed system groups to assign and ◀ move them to **Assigned managed systems**. You can

select a managed system group to see its constituent objects in the **Selected Assigned Object Members** or **Selected Available Object Members** list.

5. Save your changes and either keep the editor open with **Apply** or exit with **OK**.

Results

The situations that belong to this group are distributed to the managed systems that were assigned.

What to do next

When you view the distribution of a member situation in the Situation editor, any managed systems that were assigned to the group distribution will be disabled (dimmed). You can right-click a disabled item in the Assigned list and click Show Groups to see the situation groups that the managed system or managed system group is a member of.

Whenever you create a situation that you want to add to the situation group, click **Group**, then ⦿ **Update an existing object group** to add it.

Managed system status

The Tivoli Enterprise Portal dynamically builds the Navigator Physical view to show all the managed systems connected to the hub monitoring server and any remote monitoring servers, and their availability. Workspaces and situations at the Enterprise level keep you apprised of managed system status.

You can always see the availability of an agent in the Navigator. Offline monitoring agents are dimmed, such as the Universal Database shown here:

> 🏢 Enterprise
> > 📁 Windows Systems
> > > 🖥 MyComputer
> > > > ▪ Universal Database
> > > ◉ Windows

You can also see the status of managed systems in the Managed System Status workspace, or in a Managed System Status table view added to any workspace.

Managed System Status workspace

The Managed System Status workspace is one of the predefined workspaces for the Enterprise Navigator item, and offers a high level overview of your managed network. The workspace has a single view: *Managed System Status table view*.

Accessing the workspace

To open the Managed System Status workspace, take the following steps:

1. Click the 🏢 Enterprise Navigator item to open its default workspace.
2. Click ▦ **Workspace gallery**, then click the thumbnail graphic for **Managed System Status**.

What the workspace shows

The Managed System Status table view is a list of monitoring agents in your managed network. The Status column reports the ONLINE or OFFLINE status of the agent, with offline agents shown first.

An agent that has actually been removed from the monitored environment might appear to be offline. You can remove the display of the offline agent from the view and from the Navigator by right-clicking the table row and clicking ✎ **Clear offline entry**.

The CQ agent might be displayed in the table results. The CQ agent is a monitoring agent for the portal server and is not the actual portal server component. The status of this agent does not correlate to the status of the actual portal server. If this agent displays an offline status, there might have been a problem starting the agent or the monitoring server might not have completed refreshing.

Links to the Audit Log workspace, Agent Operations Log workspace, and Historical Export Statistics workspace

Right-click the ✎ link at the beginning of a row to select what log workspace you want to view. The log workspace is filtered to show only the messages that have been logged for the selected managed system.

Tip: If you *click* the ✎ link (instead of a right-click), the agent operations log is displayed by default. You must right-click in order to see the Audit Log or History Exports options.

Here is an example of a link from the hub monitoring server row:

▦ Managed System Status					
	Status	Name	Product	Version	Managing System	Timestamp
✎	*ONLINE	HUB_IBM-WIN1	EM	06.23.00	HUB_IBM-WIN1	01/31/13 17:00:01

This example displays a subset of the targeted Agent Operations Log linked workspace:

▦ Agent Operations Log		
Global Timestamp	Message Number	Message Text
01/31/13 17:08:29	KDSPA004	Logon validation failed. User SYSADMIN ip.pipe:#9.99.9...

This example displays a subset of the targeted Audit Log linked workspace:

▦ Audit Log					
Timestamp	Origin	Message	Result	User ID	Authorization ID
01/31/13 03:42:32	HUB_IBM-WIN1	ITM Audit services started.	0	SYSTEM	SYSTEM
01/31/13 19:46:46	HUB_IBM-WIN1	Self-describing Agent Feature enabled at the local TEMS.	0		SYSTEM

This example displays a subset of the targeted History Exports linked workspace:

▦ Historical Export Statistics					
Collection Identifer	Product Code	Attribute Group Name	Table Name	Collection Type	Collection Interval
UADVISOR_KNT_NTMEMORY	KNT	NT_Memory_64	NTMEMORY	Enterprise	00:01:00

Audit Log workspace

The Audit Log workspace displays audit events for self-describing agent related events, Warehouse Proxy Agent events, Tivoli Application Dependency Discovery Manager integration events, stop events, and start events. Audit records are stored in two formats. Audit records are stored as pure events in the ITM Audit attribute group and also in an XML formatted audit.log, see the *IBM Tivoli Monitoring Administrator's Guide* for more information.

The Audit Log workspace displays the audit records stored in the ITM Audit attribute table. You can collect data from this table historically in the IBM Tivoli Monitoring warehouse. You may also define situations to monitor the data in this table.

Tip: To quickly assess errors, look under the **Result** column for row entries with non-zero values (0 indicates a success, so any other value indicates an error). Then look at the information in the other columns to further research the error.

The audit data is normalized and based on the IBM W7 language and collected in the Audit Log table and stored in an XML format on distributed systems and in the SMF facility on z/OS. Each audit data event record answers the following questions:

Questions	Key attributes
Who: The user, application, or process that initiated the event	**Authorization ID**, **Entity**, **RunAs**, **System Name**, **User ID**
What: The type of action the event represents	See the **Operation** attributes
When: The time the event occurred	**Timestamp**
On What: The database, application, file, or permission that was manipulated	See the **Object** attributes
Where: The machine that the event occurred	See the **Origin** attributes
Where from: The system that is the source of the event	See the **Source** attributes
Where to: The targeted system of the event	See the **Target** attributes
Results: The event status and associated message	**Resource Bundle Key**, **Result**, **Message**

By default, the Audit log workspace shows audit records generated since the component was last started. In order to view audit events that occurred before the component was last restarted, you must configure the ITM Audit attribute group to collect historical data. To configure historical collection for the ITM Audit attribute complete the following steps:

1. Click ✎ History Configuration, then select **CCC Logs** under Monitored Applications.
2. Configure how long to keep Detailed data for the ITM Audit attribute group available locally.
   a. In the Select Attribute Group section, select ITM Audit.
   b. In the Configuration Controls section under Pruning select ☑ **Detailed data** and specify how long to keep the data.
   c. Click **Apply**.
3. Create a History Collection setting for the ITM Audit attribute group.

a. Right-click **CCC Logs** under Monitored Applications and select **Create new collection settings**.
   b. Supply a Name, Description and select the ITM Audit Attribute group.
   c. If you plan to historically collect monitoring server Audit records, on the Basic tab you must specify Collection Location: TEMS.
   d. On the Distribution tab, Distribute to ⦿ **Managed System (Agent)** and then select the Managed System Group or systems to start collection on.
   e. Click **OK**.

Historical Export Statistics workspace

The Historical Export Statistics workspace displays the enterprise and private historical exports to a Warehouse Proxy agent instance performed by a selected managed system. This workspace is available for monitoring servers and monitoring agents.

The statistics include the number of rows exported, the size of each row, the total amount of uncompressed data uploaded, and the timestamps for the start and end of key steps in the export process. This workspace allows you to see what historical collections are currently being exported by the managed system, determine whether or not the exports have been successful, determine the amount of time to complete an export, and determine the amount of data being exported. If the amount of exported data is excessive, you might want to reduce the collection interval for sampled attribute groups.

The historical collection statistics data is stored in the `O4SRV.KRAHIST` table. Using this table you can create custom queries and situations.

The historical export data for a managed system can be viewed using the agent's service interface. For an autonomous agent, this is the only way to view the historical export statistics data. From the service interface, select service interface request and enter the following:

```
<REPORT>
<SQLTABLE>
<TABLENAME>KRAHIST</TABLENAME>
</SQLTABLE>
</REPORT>
```

Usage notes:
* In order to view the workspace data, the selected managed system must have at minimum an IBM Tivoli Monitoring V6.3 framework or the workspace data will be empty or you might receive a `KFWITM217E Request error:` `SQL1_CreateRequest failed` message.
  For example:
  – When viewing a monitoring server's workspace data, the monitoring server must be an IBM Tivoli Monitoring V6.3 monitoring server.
  – When viewing an agent's workspace data, the agent framework used by the agent and the monitoring server it connects to and the hub monitoring server must be IBM Tivoli Monitoring V6.3.
* The workspace shows the historical collections distributed to the selected managed system, where the collection location in the historical collection is set to store at the agent. If the collection location is set to store at the managing server, the collection is displayed in the Historical Export Statistics workspace for the monitoring server the agent reports to.

- This workspace is empty if a subnode is selected. The collections are displayed in the subnode's managing system.
- The collection interval for an attribute group that contains event data is always 0 regardless of what was specified in the historical collection.
- The warehouse interval for an attribute group that contains event data is always 30 minutes regardless of what was specified in the historical collection.

Adding a Managed System Status view to a workspace

You can add a Managed System Status view to a workspace.

About this task

Use the following steps to create this table:

Procedure

1. Open the workspace where you want the view.
2. If you want the view to occupy a new space, click ▥ **Split Vertically** or ▤ **Split Horizontally** in one of the views.
3. Click ▦ **Table**.
4. Click inside the view where you want the table.
5. In the table view title bar click ✎ **Properties**.
6. In the **Query** tab, click ▯ **Click here to assign a query**. The Query editor shows a tree on the left with queries organized by product and then by attribute group:

    🗂 Queries

    ⬡ Tivoli Enterprise Monitoring Server

    🖳 Managed System

    ▯ Current Managed Systems Status

7. Under the Managed System node, Select **Current Managed Systems Status** query and click **OK**.
8. If you want *OFFLINE values to be highlighted, click the **Thresholds** tab, click inside a **Status** cell and select ***OFFLINE** from the list. Change the threshold color if you like.
9. If you want to add a title, click the **Style** tab and type new title text.
10. Click **OK** to close the Properties editor and save the table view.

EIB Change Log workspace

The Enterprise Information Base (EIB) is a database used by the Tivoli Enterprise Monitoring Server to store situation, policy, user definitions and configuration information. The EIB Change Log view is a log with details of changes made to information in the database.

Accessing the workspace

This workspace is a secondary workspace associated with the ▦ Enterprise Navigator item.

Do the following to open the EIB Change Log workspace.

1. Click the ▦ Enterprise Navigator item to open its default workspace.

2. Click ⊞ **Workspace gallery**, then click the thumbnail graphic for the **EIB Change Log** workspace. If the list is long, you might need to scroll to locate the workspace.

What the workspace shows

The EIB Change Log workspace consists of one table view (in addition to the Navigator view): EIB Change Log. The values in the log are the results of a query designed using the EIB change log attributes. The log shows, for example, the data and time at which a change was made, the ID of the user who made the last change, the name of the object (situation, policy, or other entity) that was changed, and the type of operation that took place.

Self-Monitoring Topology workspace

The Self-Monitoring Topology workspace provides a high level overview of your managed infrastructure and its health and introduces you to the Tivoli Enterprise Portal self-monitoring capabilities.

Accessing the workspace

The Self-Monitoring Topology workspace is a secondary workspace associated with the ⊞ Enterprise Navigator item.

Do the following to open this workspace.
1. Click the ⊞ Enterprise Navigator item to open its default workspace.
2. Click ⊞ **Workspace gallery**, then click the thumbnail graphic for **Self-Monitoring Topology**. If the list is long, you might need to scroll to locate the workspace.

What the workspace shows

The workspace contains three views:

TMS Infrastructure - Base

> TMS Infrastructure - Base is a topological display of the Tivoli Management Services (TMS) components and their availability status. The Tivoli Enterprise Portal Server connects to the hub Tivoli Enterprise Monitoring Server, as does the Tivoli Data Warehouse, the Summarization and Pruning agent, and any remote monitoring servers. See TMS Infrastructure view for a description of the view features and tools for manipulating the topology.

Users Logged on to TEPS

> Users Logged on to TEPS is a table view of the users logged on to the portal server, with the IP address of the client computer and logon time stamp.

Managed Systems count per Managing System

> This view is a bar chart showing a count of managed systems grouped by *managing system*. Also called *through node*, this is the parent system that collects data from the agent or subagent. If the agent is directly connected to the hub Tivoli Enterprise Monitoring Server, the hub monitoring server is the managing system; if the agent is connected to a remote monitoring server, then that is the managing system. An agent with subagents is the managing system for each subagent.

Tivoli Enterprise Monitoring Server status

For self-monitoring, the Tivoli Enterprise Monitoring Server has attributes that are on display in workspaces and used in situations that are available through the Enterprise Navigator item.

Manage Tivoli Enterprise Monitoring Servers workspaces

Use the Manage Tivoli Enterprise Monitoring Servers workspaces and situations to gather information for diagnosing and correcting typical monitoring server configuration issues.

The Manage Tivoli Enterprise Monitoring Servers workspace is available with the other workspaces from the Enterprise Navigator item. Open the workspace as a quick health check whenever you have updated any components. While an ▦ Enterprise-level workspace is open, click ▦ **Workspace Gallery**, then scroll to and click the **Manage Tivoli Enterprise Monitoring Servers** workspace.

The **Tivoli Enterprise Monitoring Server Topology** view shows the hub monitoring server and all remote monitoring servers that connect to it in the enterprise, with a status indicator overlaying the 🖳 monitoring server icon. Move the mouse pointer over the icon to see an information box with identifying data for the monitoring server, such as the host address and the status. Right-click a monitoring server icon to see and select a workspace to open with more information about that monitoring server.

The **Monitoring Servers** view lists each Tivoli Enterprise Monitoring Server and whether it is online or offline. The hub is distinguished by an 🖳 icon after the name. This view persists in all the Manage Tivoli Enterprise Monitoring Servers workspaces and in the same position.

Click a 🖉 link to open one of the detailed workspaces with data from the monitoring server:

Manage Tivoli Enterprise Monitoring Servers - Enterprise Catalogs
> This workspace is displayed after you select the **Installed Catalogs - Enterprise View** link from the hub monitoring server. The link to this workspace is available only for a hub monitoring server.
>
> The application catalogs that are installed on every remote monitoring server in the enterprise are compared with those on the hub monitoring server and any discrepancies are highlighted in the **Status** column. The timestamp was taken when the configuration file was last edited.
>
> > 🟩 **OK** indicates that the timestamps and column counts of the hub and remote monitoring servers match.
> >
> > 🟦 **Timestamp Mismatch** indicates that the time the remote monitoring server was last updated does not match the time the hub was last updated.
> >
> > 🟨 **Hub Newer** indicates that the hub monitoring server was updated more recently than a remote monitoring server.
> >
> > 🟧 **Missing At RTEMS** indicates that the catalog ID that is present at the hub monitoring server is missing at a remote monitoring server.
> >
> > 🟥 **Count Mismatch** indicates that the column count at the remote monitoring server does not match the column count at the hub monitoring server.

■ **RTEMS Newer** indicates that a remote monitoring server was updated more recently than the hub monitoring server.

■ **Missing At Hub** indicates that the catalog ID that is present on a remote monitoring server is missing at the hub monitoring server.

Typically, the hub and remote monitoring servers have similar or an identical set of catalogs. If any of the following catalogs are missing at the hub, it is most likely not a problem (especially if this is a distributed hub): PDSSTATS, OMSMS, KSD, and MV.

Manage Tivoli Enterprise Monitoring Servers - Installed Catalogs
This workspace is displayed after you select the **Installed Catalogs** link from a remote monitoring server. The Remote Server shows a comparison of the application catalogs that are installed on the selected Remote Tivoli Enterprise Monitoring Server with those on the hub monitoring server. Here you can get a report of catalog out-of-sync conditions. The "Status" column has the same values as those shown in the Manage Tivoli Enterprise Monitoring Servers - Enterprise Catalogs workspace.

The link to this workspace is available to all remote monitoring servers, but not the hub.

Manage Tivoli Enterprise Monitoring Servers - Protocols
This workspace is displayed after you select the **Protocols** link. The workspace shows the Local and Global Communications Protocol (location broker) entries for the monitoring server. These are the definitions that were created during installation for connection between components, such as the protocol, IP address, and port number. The global communications protocol entries are available only for the hub monitoring server. This workspace is available only for monitoring servers that are at Version 6.2.2 or later.

Manage Tivoli Enterprise Monitoring Servers - Situation Status
This workspace is displayed after you select the **Situation Status** link. The workspace reports the situations that are started on the monitoring server and has summarized and detailed views of the situations that count the number of times the status type has persisted on the monitoring server.

Manage Tivoli Enterprise Monitoring Servers - System Information
This workspace is displayed after you select the **System Information** link. The workspace supplies detailed information about the monitoring server process and system configuration. It is available for monitoring servers from Version 6.2.2 onward. The Server Configuration view shows the parameters that have been configured for the monitoring server.

These subordinate workspaces are available only through their parent Manage Tivoli Enterprise Monitoring Servers workspace. (The exception is when you are in workspace administration mode, although no data is returned if you access the workspaces directly.)

Manage Tivoli Enterprise Monitoring Server situations

The Manage Tivoli Enterprise Monitoring Server predefined situations are used to monitor the status of the monitoring server catalogs.

These situations monitor the status of the Tivoli Enterprise Monitoring Server catalogs:

Hub_Catalog_Critical
This situation becomes true and an event is opened when an application

catalog that is installed at a remote monitoring server is *not* installed at the
hub monitoring server. This can happen when application support has not
been selected for the hub during installation. Default: Do not □ **Run at
startup**. ⚙ Show Formula:

```
( Status == Missing At Hub) AND (Catalog != 'PDSSTATS')
```

Hub_Catalog_Warning
This situation becomes true and an event is opened when an application
catalog that is installed at a remote monitoring server is newer than what
is installed at the hub monitoring server. This is a warning that the
application support for the component is out of date. This might not be a
problem if the hub has not been upgraded yet. Default: Do not □ **Run at
startup**. ⚙ Show Formula:

```
( Status == RTEMS Newer)
```

Remote_Catalog_Warning
This situation becomes true and an event is opened when an application
catalog that is installed at the hub monitoring server is not also installed at
a remote monitoring server; or an application catalog that is installed at the
hub is newer than what is installed at a remote monitoring server.

This might not be a problem if the remote monitoring server has not been
upgraded yet. Also, the remote system might be intentionally not
supporting the subject application, or might only be supporting an older
version of the application. Default: Do not □ **Run at startup**. ⚙ Show
Formula:

```
((Status == Missing At Hub) AND (Catalog != 'PDSSTATS')) OR
(Status == Hub_Newer)
```

Managing monitoring agents

With the Tivoli Enterprise Portal you can deploy and maintain distributed Tivoli
Enterprise Monitoring Agents, such as the Sybase Server monitoring agent, and
connect to a distributed Hub Tivoli Enterprise Monitoring Server.

Before you can remotely install and configure distributed monitoring agents on a
computer, the computer must have an OS enterprise monitoring agent installed for
the operating system the monitoring product will run under. When the OS agents
have been installed, the Navigator Physical view adds an item for each online
managed system. Monitoring agents that do not support the remote agent
deployment feature will not display the **Add Managed System**, **Configure**, and
Remove options in the Navigator pop-up menu. The types of managed systems
that you can add to a computer depend on what agent bundles are in the depot on
the monitoring server to which the OS agent is connected. See the *IBM Tivoli
Monitoring Installation and Setup Guide* for more information.

Agent deployment workspaces

Use the agent deployment workspaces to see your package list, the results of a
deployment prerequisite check, and the progress of the agent deployment process.

Deploy Depot Package List workspace

Use the Deploy Depot Package List workspace to see the installation packages that
are available in the agent depot at the Tivoli Enterprise Monitoring Server.

This is one of the predefined workspaces available from the Enterprise Navigator
item.

Accessing the workspace

1. Click the ▦ Enterprise item in the Navigator Physical view.
2. Click ▦ **Workspace gallery**.
3. Scroll to **Deploy Depot Package List**, then click the thumbnail graphic to open the workspace.

What the workspace shows

The Deploy Depot Package List workspace shows a table listing all of the agent installation packages available in the agent depot for remote deployment. For each package, the table includes the product ID, version, host type, and other details about the package.

For more information about the agent depot, see "Deploying monitoring agents across your environment" in the *IBM Tivoli Monitoring Installation and Setup Guide* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/ com.ibm.itm.doc_6.3fp2/install/itm_install.htm).

Related reference:

"Deployment Status Summary by Transaction workspace" on page 402
The Deployment Status Summary by Transaction workspace shows summary status information about pending and completed asynchronous remote agent deployments, sorted by transaction. This is one of the predefined workspaces available from the Enterprise Navigator item.

Deployment Status Summary workspace

Use the Deployment Status Summary workspace to get ongoing status detail about pending and completed asynchronous remote agent deployments. The charts and tables have visual indicators for quick assessment of the success rate and any problems that arise.

Accessing the workspace

1. Click the ▦ Enterprise item in the Navigator Physical view.
2. Click ▦ **Workspace gallery**.
3. Scroll to **Deployment Status Summary**, then click the thumbnail graphic to open the workspace.

What the workspace shows

Deployment Status Summary
    Deployment Status Summary is a table view showing total counts of the pending asynchronous agent deployments. For each target managed system, the table shows the number of deployments in each status (succeeded, failed, retrying, queued, and in progress).

Status Summary of Complete Deployments
    Status Summary of Complete Deployments is a stacked bar chart showing the asynchronous agent deployments that have finished processing. For each target managed system, the chart shows how many deployments succeeded and how many failed.

Status Summary of Incomplete Deployments
    Status Summary of Incomplete Deployments is a stacked bar chart showing the asynchronous agent deployments that have not yet finished processing. For each target managed system, the chart shows the deployments that are queued, in progress, or retrying.

Deployment Status Detail

Deployment Status Detail is a table view showing detailed information about all asynchronous agent deployments. The table shows the status, product ID, target managed system, and other details about each deployment request. Click 🖉 on a row to link to a more detailed workspace for the selected deploy group or product. See "Deployment Status by Deploy Group workspace" on page 404 and "Deployment Status by Product workspace" on page 405 for details.

Tip: The timestamps for a transaction are generated by the monitoring server that is processing the deployment. The hub monitoring server inserts the transaction first. You might have a start time that is earlier than the latest timestamp because the remote monitoring server is behind the hub monitoring server. To avoid this confusion, you might want to sync the time on the hub monitoring server and remote monitoring server.

Deployment notes

Asynchronous remote agent deployment

Agents can be remotely deployed asynchronously as well as synchronously. With asynchronous deployment, you can request that a second remote agent be deployed, even if the previously deployed agent has not initialized fully. You can also group agents for bulk remote deployment.

See "Deploying monitoring agents across your environment" in the *IBM Tivoli Monitoring Installation and Setup Guide* (http://pic.dhe.ibm.com/ infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/install/ itm_install.htm) for details.

CLI commands for agent deployment

The `tacmd getdeploystatus` command reports the deployment status from the command line.

The `tacmd clearDeployStatus` command can be used to clear all or a filtered list of entries in the status table.

The `tacmd checkprereq` command checks for prerequisites on the target system before you begin deploying a monitoring agent there. You can review the results in the Deployment Status Summary workspace or with the `tacmd getDeployStatus` command.

The `tacmd restartFailedDeploy` command restarts failed deployments. You can restart all the failed entries in the status table or filter the table entries

The *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/ infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/ic/ landing_cmdref.htm) describes these commands and their syntax.

Deployment Status Summary by Transaction workspace

The Deployment Status Summary by Transaction workspace shows summary status information about pending and completed asynchronous remote agent deployments, sorted by transaction. This is one of the predefined workspaces available from the Enterprise Navigator item.

Accessing the workspace

1. Click the ▦ Enterprise item in the Navigator Physical view.
2. Click ▦ **Workspace gallery**.

3. Scroll to **Deployment Status Summary by Transaction**, then click the thumbnail graphic to open the workspace.

What the workspace shows

The workspace contains four views:

Deployment Status Summary by Transaction
Deployment Status Summary by Transaction is a table view showing total counts of the pending asynchronous agent deployment transactions. For each target managed system, the table shows the number of deployment transactions in each status (succeeded, failed, retrying, queued, and in progress).

Status Summary of Complete Deployments by Transaction
Status Summary of Complete Deployments by Transaction is a stacked bar chart showing the asynchronous agent deployment transactions that have finished processing. For each target managed system, the chart shows how many succeeded and how many failed.

Status Summary of Incomplete Deployments by Transaction
Status Summary of Incomplete Deployments by Transaction is a stacked bar chart showing the asynchronous agent deployment transactions that have not yet finished processing. For each target managed system, the chart shows the transactions that are queued, in progress, or retrying.

Deployment Status Detail
Deployment Status Detail is a table view showing detailed information about all asynchronous agent deployments. The table shows the status, product ID, target managed system, and other details about each deployment request. Click 🖉 on a row to link to a more detailed workspace for the selected deploy group or product. See "Deployment Status by Deploy Group workspace" on page 404 and "Deployment Status by Product workspace" on page 405 for details.

Tip: The timestamps for a transaction are generated by the monitoring server that is processing the deployment. The hub monitoring server inserts the transaction first. You might have a start time that is earlier than the latest timestamp because the remote monitoring server is behind the hub monitoring server. To avoid this confusion, you might want to sync the time on the hub monitoring server and remote monitoring server.

Deployment notes

Asynchronous remote agent deployment
Agents can be remotely deployed asynchronously as well as synchronously. With asynchronous deployment, you can request that a second remote agent be deployed, even if the previously deployed agent has not initialized fully. You can also group agents for bulk remote deployment.

See "Deploying monitoring agents across your environment" in the *IBM Tivoli Monitoring Installation and Setup Guide* (http://pic.dhe.ibm.com/ infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/install/ itm_install.htm) for details.

CLI commands for agent deployment
The `tacmd getdeploystatus` command reports the deployment status from the command line.

The **tacmd clearDeployStatus** command can be used to clear all or a filtered list of entries in the status table.

The **tacmd checkprereq** command checks for prerequisites on the target system before you begin deploying a monitoring agent there. You can review the results in the Deployment Status Summary workspace or with the **tacmd getDeployStatus** command.

The **tacmd restartFailedDeploy** command restarts failed deployments. You can restart all the failed entries in the status table or filter the table entries

The *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/ infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/ic/ landing_cmdref.htm) describes these commands and their syntax.

Related reference:
"Deploy Depot Package List workspace" on page 400
Use the Deploy Depot Package List workspace to see the installation packages that are available in the agent depot at the Tivoli Enterprise Monitoring Server.

Deployment Status by Deploy Group workspace

The Deployment Status by Deploy Group workspace shows status information about asynchronous remote agent deployments for a particular deploy group.

This workspace is accessed from the Deployment Status Summary workspaces.

Accessing the workspace

1. Open the Deployment Status Summary workspace or Deployment Status Summary by Transaction workspace from the ▦ Enterprise Navigator item.
2. In the Deployment Status Detail table, click the 🖉 link in the row of the deploy group for which you want detailed status.
3. Click **Deployment Status by Deploy Group** to open the workspace.

What the workspace shows

The workspace contains four views:

Deployment Status Summary by Deploy Group

> Deployment Status Summary is a table view showing total counts of the pending asynchronous agent deployments for the selected deploy group. For each target managed system, the table shows the number of deployments in each status (succeeded, failed, retrying, queued, and in progress).

Status Summary of Complete Deployments
> Status Summary of Complete Deployments is a stacked bar chart showing the asynchronous agent deployments for the selected deploy group that have finished processing. For each target managed system, the chart shows how many deployments succeeded and how many failed.

Status Summary of Incomplete Deployments
> Status Summary of Incomplete Deployments is a stacked bar chart showing the asynchronous agent deployments for the selected deploy group that have not yet finished processing. For each target managed system, the chart shows the deployments that are queued, in progress, or retrying.

Deployment Status Detail by Deploy Group
> Deployment Status Detail is a table view showing detailed information

about all asynchronous agent deployments for the selected deploy group. The table shows the status, product ID, target managed system, and other details about each deployment request.

Deployment Status by Product workspace

The Deployment Status by Product workspace shows status information about asynchronous remote agent deployments for a particular product ID.

This workspace is accessed from the Deployment Status Summary workspaces.

Accessing the workspace

1. Open the Deployment Status Summary workspace or Deployment Status Summary by Transaction workspace from the ▦ Enterprise Navigator item.
2. In the Deployment Status Detail table, click the🖉 link in the row for which you want detailed status by product.
3. Click **Deployment Status by Product** to open the workspace.

What the workspace shows

The workspace contains four views:

Deployment Status Summary by Product

> Deployment Status Summary is a table view showing total counts of the pending asynchronous agent deployments for the selected product ID. For each target managed system, the table shows the number of deployments in each status (succeeded, failed, retrying, queued, and in progress).

Status Summary of Complete Deployments
> Status Summary of Complete Deployments is a stacked bar chart showing the asynchronous agent deployments for the selected product ID that have finished processing. For each target managed system, the chart shows how many deployments succeeded and how many failed.

Status Summary of Incomplete Deployments
> Status Summary of Incomplete Deployments is a stacked bar chart showing the asynchronous agent deployments for the selected product ID that have not yet finished processing. For each target managed system, the chart shows the deployments that are queued, in progress, or retrying.

Deployment Status Detail by Product
> Deployment Status Detail is a table view showing detailed information about all asynchronous agent deployments for the selected product ID. The table shows the status, target managed system, and other details about each deployment request.

Stopping a monitoring agent

You can stop running any of the distributed monitoring agents other than the OS agents, which are used for agent deployment and configuration.

About this task

To stop a monitoring agent:

Your user ID must have Agent Management Start/Stop permission to use this feature.

Procedure

1. In the Navigator physical view, right-click the ⬚ agent-level Navigator item.
2. Click ⬤ **Stop**. The agent goes offline and the Navigator item is dimmed. The agent will not come online again until it has been restarted on the monitoring server to which it is assigned.

Recycling a monitoring agent

You can recycle a monitoring agent.

About this task

To recycle a monitoring agent:

👤 Your user ID must have Agent Management Start/Stop permission to use this feature.

Procedure

1. In the Navigator physical view, right-click the ⬚ agent-level Navigator item.
2. Click ⬚ **Restart** to stop, then start the monitoring agent.

Starting an offline monitoring agent

An offline monitoring agent can be started.

About this task

To start an offline monitoring agent:

👤 Your user ID must have Agent Management Start/Stop permission to use this feature.

Procedure

1. In the Navigator physical view, right-click the offline ⬚ agent-level Navigator item.
2. Click ▷ **Start**. The request to start the monitoring agent is sent to the monitoring server to which it is connected. Depending on your monitoring configuration, it might take a few moments before the agent starts running and you see the Navigator item enabled.

Results

Note: If the monitoring agent does not start and you get an error message, the computer might be unavailable.

Adding a managed system through the Tivoli Enterprise Portal

You can install and configure managed systems where OS monitoring agents are installed and running.

Before you begin

To use this feature, your user ID must have Manage permission for Agent Management.

Procedure

1. In the Navigator physical view, right-click the 📁 system-level Navigator item for the computer where you want to install the monitoring agent. In this example, the computers named ORANGE, PEAR, CABBAGE, and ONION are available.

      ▦ Enterprise

            🖥 Linux Systems

                📁 ORANGE

                📁 PEAR

            🖥 Windows Systems

                📁 CABBAGE

                📁 ONION

2. Click ➕ **Add Managed System** to open the Select a Monitoring Agent window. The agents shown in this list are those available for the operating system on which this computer runs. The two-digit version number is followed by a two-digit release number and a modification number of up to five digits.

3. Highlight the type of the monitoring agent to install and click **OK**. For some agent types, the new managed system operation is queued and the transaction ID is displayed. For other agent types, a wizard is provided for you to configure the agent on this system.

4. Complete the fields to configure the agent, clicking **Next** and **Back** to move among the pages.

5. On the **Agent** page, establish the operating system user ID under which the agent will run on the managed system. *Windows*: Either accept the default to start the managed system with your user ID (you can also select the check box to **Allow service to interact with desktop** to enable remote control) or select **Use this account** and fill in the user name and password under which the agent will run.

   Non-Windows: Enter the **Username** under which the agent will run and the **Group name**.

6. Click **Finish** to complete the managed system configuration. If any of the information provided is invalid, you will receive an error message and be returned to the configuration window. Check your entries and edit as appropriate to configure correctly. Installation and setup begins and might take several minutes to complete depending on your Tivoli monitoring configuration, the location of the managed system, and the type of monitoring agent.

7. After the managed system has been added to the enterprise, click 🔗 **Apply Pending Updates** in the Navigator view toolbar. The new managed system (such as 🔲 Universal Database) is displayed below the system Navigator item.

Configuring a managed system

After adding a managed system to your enterprise, you can open the Configure Managed System window whenever you want to adjust the configuration settings.

Before you begin

To use this feature, your user ID must have Manage permission for Agent Management.

Procedure

1. Right-click the ⊙ Navigator item for the agent to configure or upgrade.
2. Click 🔑 **Configure** to open the Configure Managed System window.
3. Edit the fields to configure the agent, clicking **Next** and **Back** to move among the pages. Any pages besides **Agent** are specific to the agent type.
   - Performance Analyzer, Summarization and Pruning Agent, and Warehouse Proxy: See "Tivoli Data Warehouse solutions" in the *IBM Tivoli Monitoring Installation and Setup Guide*.
   - Non-base agents: See your product's installation guide in the IBM Tivoli Monitoring Information Center (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/welcome.htm) or on IBM Tivoli Documentation Central (https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Documentation%20Central).
4. On the **Agent** page, establish the user ID that will be used to maintain the agent:

   *Windows***:**
   > Accept the default ⊙ **Use local system account** to use your Tivoli Enterprise Portal user ID. You can also select ☐ **Allow service to interact with desktop** to enable remote control. Or select ○ **Use this account** and fill in the user name and password under which the agent will be controlled.

   *Non-Windows***:**
   > Enter the **Username** under which the agent will run and the **Group name**.

5. Click **Finish** to complete the managed system configuration. If any of the information provided is invalid, you will receive an error message and be returned to the configuration window. Check your entries and edit as appropriate to configure correctly.

Updating a monitoring agent

When a new version of a distributed monitoring agent is released, you can apply the new version locally or remotely to one managed system at a time, or to many simultaneously. Use the Configure Managed System window to apply the update.

Note: Agent updates are automatic and this procedure is not required if your monitoring agent is at Version 6.2.3 or higher unless the self-describing capability is disabled.

Before you begin

A patch is for a specific version of the agent. The prereq field of the depot entry contains an agent of the same type as the agent itself. For example, if 6.1 UD (DB2 monitoring) is the selected agent, the prereq field in the depot entry must contain a deployment bundle notation such as ud:061000000.

You cannot upgrade an agent to a new version with this procedure, but rather with the **tacmd upgradeAgent** command, which is described in the *IBM Tivoli Monitoring Command Reference* (http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/topic/com.ibm.itm.doc_6.3fp2/ic/landing_cmdref.htm). Do not try to change the configuration of a managed system at this time because your changes will be lost when the installation of the updates begins.

To use this feature, your user ID must have Manage permission for Agent Management.

About this task

Complete these steps to apply a patch for a monitoring agent:

Procedure

1. Right-click the ⊙ Navigator item for the agent that you want to upgrade.
2. Click ⚲ **Configure** to open the Configure Managed System window.
3. Click the **Agent** tab.
4. Compare the installed version of the monitoring agent with any available product updates, then highlight the row of the agent to update and click **Install Updates**.

Results

Installation of the updates begins and might take several minutes to complete. The list that displays reflects the contents of the deployment depot. If **Install Updates** is disabled, one or more of the following conditions exist:

- The depot entry does not match the product type.
- The **VVRR** fields for the agent and the depot entry are *not* the same, where VV is the version number and RR is the revision number. For example, an entry of `0610` prevents you from applying a fix pack intended for a version 6.2 agent.
- The depot entry is at an older version than the agent.
- The host version field of the depot entry does not contain the host platform for the agent.
- The prereq field of the depot entry does not contain an agent of the same type as the agent itself. For example, if 6.1 UD (DB2 monitoring) is the selected agent, the prereq field in the depot entry must contain a deployment bundle notation such as ud:061000000, which is one way to denote a patch deployment bundle.

Removing a managed system and monitoring agent

Use ☒ **Remove** to stop a managed system and permanently remove it from the Navigator Physical view.

(You can add it back in at a later time using ⊞ **Add Managed System**.)

Before you begin

To use this feature, your user ID must have Manage permission for Agent Management.

About this task

After you have removed all managed systems, you can choose to completely uninstall the monitoring agent. If you plan to reinstall the monitoring agent in future, it is best to leave it installed.

Procedure

1. Right-click the ⊙ agent Navigator item.
2. Click ☒ **Remove**.

3. Click **Yes** when a message asks you to confirm the removal of the managed system; or **No** to cancel.

4. If another message asks you to confirm that you want to uninstall the monitoring agent, click **Yes** to permanently uninstall the monitoring agent from the computer; or **No** to cancel.

Results

The agent item no longer appears in the Navigator.

Related tasks:

"Removing an offline managed system Navigator item" on page 32
Use the Managed System Status workspace to clear offline managed systems from the Navigator view.

Optimizing IBM Tivoli Monitoring resources

Here are some tips for making the most of your monitoring resources:

- Tivoli Enterprise Portal must run every query you use for different chart and table views. If you use the same query for these views, data retrieval speeds up and you conserve Tivoli Enterprise Portal Server and monitoring server resources.

- Tivoli Enterprise Portal uses queries to specify the attributes to retrieve for a table or chart. After that data is retrieved, you can use the Properties editor **Filters** tab to filter out any unwanted data from the view. You can use the same filtering capability in the query itself to filter the data *before* it is retrieved. This *pre-filtering* means the portal server has less data to retrieve and therefore speeds up the process.

- You can control the time span shown in a chart or table by adding a filter (through the Properties window) or by using ◢ **time span** in the chart or table view toolbar. Using the **Filters** tab causes the filtering to be done *after* the data has been received at the client. ◢ **time span** pre-filters at the server, so only the rows that pass the test are returned, and you conserve network resources.

- When you create a situation, the interval is set to 15 minutes by default. Set the situation to a longer interval between data samplings to reduce traffic on the monitoring server. This is done in the **Sampling interval area** of the **Formula** tab in the Situation editor.

- A workspace gathers new data samplings when you first open it. You can set the workspace to refresh automatically anywhere from every 30 seconds to once an hour. If you have workspaces set for automatic refresh, set the interval as high as possible to reduce network traffic (requests from the portal client go through the portal server to the monitoring server, then to the monitoring agent, and the information is sent back through the same path).

- Event data is stored in the KFW tables on the portal server. To ensure the performance of your portal server, you must periodically prune this data. You can do this by controlling how long the event data is kept (using the KFW_EVENT_RETENTION parameter), at what time of the day the pruning

begins (using the KFW_PRUNE_START parameter), and when the pruning ends (using the KFW_PRUNE_END parameter). Set these parameters in the Tivoli Enterprise Portal Server environment file (KFWENV), as described in the *IBM Tivoli Monitoring Administrator's Guide*.

• A poorly written situation can have a surprisingly large impact on CPU usage on a monitoring server. Using wildcard characters (*) in situations is one of the most common mistakes made in writing situations. Using fields longer than required can also cause additional work to evaluate the situation. Making the changes below have, in some instances, reduced CPU usage on a monitoring server from close to full utilization to nearly completely idle. The CPU usage manifests itself in the kdsmain process. See "High CPU usage on the monitoring server" on page 325 for more information.

Organization of predefined workspaces

The Enterprise Navigator item has workspaces that query the Tivoli Enterprise Monitoring Server. Use these predefined workspaces to get status information about the monitoring server and monitoring agents and about situations and policies.

Use the ⊞ **Workspace Gallery** to see what is available for the Navigator item. These are the Enterprise Navigator item workspaces and the workspaces they link to:

Enterprise Status
> The default workspace is Enterprise Status, which gives an overview of the situation event status throughout your enterprise.
> 🔲 Event Details - Similar by Situation Name
> 🔲 Event Details - Similar by Source
> 🔲 Event Details - Similar by Resource

Manage Tivoli Enterprise Monitoring Servers
> The Manage Tivoli Enterprise Monitoring Servers workspaces provide a visual health check of the monitoring servers in your enterprise and the application support that has been applied.
> 🔲 Installed Catalogs - Enterprise View
> 🔲 Installed Catalogs - Remote Server
> 🔲 Protocols
> 🔲 Situation Status
> 🔲 System Information

Managed System Status
> The Managed System Status is a list of monitoring agents in your managed network and their ONLINE or OFFLINE status. The linked workspaces are only available for online managed systems.
> 🔲 Audit Log
> 🔲 Agent Operations Log
> 🔲 History Exports

EIB Change Log
> This workspaces displays entries in the Enterprise Information Base (EIB) log. The EIB is a database used by the Tivoli Enterprise Monitoring Server to store situation, policy, user definitions, and configuration information.

Self-Monitoring Topology
> The Self-Monitoring Topology workspace provides a high level overview of your managed infrastructure and its health.

Deploy Depot Package List

The Deploy Depot Package List workspace shows the installation packages that are available in the agent depot.

Deployment Status Summary

The Deployment Status Summary workspace shows summary status information about remote agent deployments.

- Deployment Status by Deploy Group
- Deployment Status by Product

Deployment Status Summary by Transaction

The Deployment Status Summary by Transaction workspace shows summary status information about remote agent deployments, sorted by transaction.

Related reference:

"Enterprise Status workspace" on page 57
At the highest level of the Navigator Physical view is the ▦ Enterprise item. The default workspace for the Enterprise is called Enterprise Status.

"Manage Tivoli Enterprise Monitoring Servers workspaces" on page 398
Use the Manage Tivoli Enterprise Monitoring Servers workspaces and situations to gather information for diagnosing and correcting typical monitoring server configuration issues.

"Managed System Status workspace" on page 392
The Managed System Status workspace is one of the predefined workspaces for the Enterprise Navigator item, and offers a high level overview of your managed network. The workspace has a single view: *Managed System Status table view*.

"EIB Change Log workspace" on page 396
The Enterprise Information Base (EIB) is a database used by the Tivoli Enterprise Monitoring Server to store situation, policy, user definitions and configuration information. The EIB Change Log view is a log with details of changes made to information in the database.

"Self-Monitoring Topology workspace" on page 397
The Self-Monitoring Topology workspace provides a high level overview of your managed infrastructure and its health and introduces you to the Tivoli Enterprise Portal self-monitoring capabilities.

"Agent deployment workspaces" on page 400
Use the agent deployment workspaces to see your package list, the results of a deployment prerequisite check, and the progress of the agent deployment process.

Tivoli Enterprise Monitoring Server situations

The Enterprise Navigator item has a set of predefined situations. They appear within the Tivoli Enterprise Monitoring Server entry of the Situation editor.

Deploy_Failed

This situation becomes true and an event is opened when one or more agent deployments has failed. Default: Do not ☐ **Run at startup**. ⚛ Show Formula:

(Status == Failed)

Deploy_Retrying

This situation becomes true and an event is opened when one or more agent deployments is in the FailedRetry state, which means it has failed an attempt to deploy and is making another attempt. Default: Do not ☐ **Run at startup**. ⚛ Show Formula:

(Status == 'Failed Retrying')

Hub_Catalog_Critical

This situation becomes true and an event is opened when an application catalog that is installed at a remote monitoring server is *not* installed at the hub monitoring server. This can happen when application support has not been selected for the hub during installation. Default: Do not □ **Run at startup**. Show Formula:

```
( Status == Missing At Hub) AND (Catalog != 'PDSSTATS')
```

Hub_Catalog_Warning

This situation becomes true and an event is opened when an application catalog that is installed at a remote monitoring server is newer than what is installed at the hub monitoring server. This is a warning that the application support for the component is out of date. This might not be a problem if the hub has not been upgraded yet. Default: Do not □ **Run at startup**. Show Formula:

```
( Status == RTEMS Newer)
```

Remote_Catalog_Warning

This situation becomes true and an event is opened when an application catalog that is installed at the hub monitoring server is not also installed at a remote monitoring server; or an application catalog that is installed at the hub is newer than what is installed at a remote monitoring server.

This might not be a problem if the remote monitoring server has not been upgraded yet. Also, the remote system might be intentionally not supporting the subject application, or might only be supporting an older version of the application. Default: Do not □ **Run at startup**. Show Formula:

```
((Status == Missing At Hub) AND (Catalog != 'PDSSTATS')) OR
(Status == Hub_Newer)
```

TEMS_Busy

This situation becomes true and an event is opened when the Tivoli Enterprise Monitoring Server is continually busy. Default: Do not □ **Run at startup**. Show Formula:

```
( Product == 'EM' AND Reason == 'H')
```

Appendix A. Formula functions

Use the functions provided in these editors to write formulas: Situation, Query, Properties Filters tab, and Properties Thresholds tab.

Formula function types

With over a dozen cell and group functions available, you can build sophisticated filters for reporting only what matters to you, thus saving time and network resources. Likewise, you can build thresholds in table views and conditions in situations that trigger alerts to keep you informed of important changes in your managed network.

The functions that are available depend on the type of attribute (numeric, text, or time) and in which editor you are working: Query, Situation, Properties (Filters and Thresholds tabs), or History Collection Configuration (Filter tab).

Cell functions

This table maps the 123 numeric, ABC text, and ⏱ timestamp attribute types to the cell and group formula functions and to the editors where they are available.

Table 10. Formula functions and the editors they are available in

Cell function	Query specification	View properties: Filters; Thresholds	Situation: Formula	History Collection Configuration: Filter
▦ BETWEEN	123	123 (Thresholds only)	Replicate the ▦ **See if a value is within a range** functionality by repeating the attribute in the same row and using the > or >= and < or <= operators.	
△ CHANGE (Formula editor row 1 only)			123	
▦ DATE	⏱	⏱	⏱	⏱
∅ ISNULL		123 ABC ⏱		
🔍 MISSING (Formula editor row 1 only)			ABC	
% PCTCHANGE (Formula editor row 1 only)			123	
⧉ REGEX		ABC ⏱		
abc SCAN	ABC ⏱	ABC ⏱	ABC ⏱	ABC ⏱
abc STR	ABC ⏱	ABC ⏱	ABC ⏱	ABC ⏱
◷ TIME		⏱	⏱	
⏱ TIMESPAN		⏱		
ᵥ VALUE	123 ABC ⏱	123 ABC ⏱	123 ABC ⏱	123 ABC ⏱

Table 10. Formula functions and the editors they are available in (continued)

Cell function	Query specification	View properties: Filters; Thresholds	Situation: Formula	History Collection Configuration: Filter
VALUE IN (IN is a relational operator for the Value function)		ABC ⊗	ABC	ABC

Where you see the notation, "Formula editor row 1 only", the function must be used in the first-row cell of the formula editor; subsequent rows will be disabled. Additionally, a formula with this function cannot include any other functions that have this restriction or be combined with any group functions.

Group functions

Use the group functions in the Query editor or Situation editor for attribute groups that return multiple rows of data or a single-row attribute group that has been configured for historical data collection.

With the exception of COUNT, which can be used by all attribute types, group functions are available only for numeric attributes.

Table and chart views require that a time range be set to show a span of data samplings.

Group function	Query Specification	Situation Formula
AVG	fx	123
COUNT	fx	123 ABC ⊗
MAX	fx	123
MIN	fx	123
Σ SUM	fx	123

Note:

1. Time Span selection is not available for a table or chart view if it uses a query that includes a group function.
2. When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x456c for hex value 456c.
3. If the expression that is responsible for a situation event uses a group function, the situation event results workspace shows neither the attribute value threshold color nor the expression that should display when you hover over an initial situation value.

Examples

After you open a workspace that has query-based views with filters applied or after you start a situation, the sampled value of each referenced attribute is compared with the expression. A true comparison causes the value to be filtered into the view or to open an event.

For example, the following expression uses the VALUE formula function and can be read as, "When the VALUE of Day of Week is greater than or equal to 2 (Monday) and the VALUE of Day of Week is less than or equal to 6 (Friday), the comparison is true."

	Day of Week	Day of Week
1	>= 2	<= 6

Here is another example using the MIN formula function, which is a *group function*. This formula finds the process with the lowest priority base. If this process also uses 20% or more privileged time, the comparison is true.

	Priority Base	% Privileged Time
1	▢ == *TRUE	>= 20

Attribute and attribute group characteristics

Formula functions treat attributes as numeric, text or timestamp. Attributes also have other characteristics that further categorize them and inform your choice of formula function or query-based view..

Numeric

Numeric attributes can represent a percentage, seconds, bytes or some other unit of measurement.

When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x123ab for hex value 123ab. In z/OS applications, for example, DASD volume serial numbers are hexadecimal.

Text   Single quotes are required around multiple-word text values. The timestamp and, usually, enumerated attributes are treated as text attributes.

Words with apostrophes, such as **that's** or **agent's**, cannot be used. Instead, consider using a wildcard or, if appropriate, use **Return a subset of the string or** or **Scan for string within a string** and specify the letters preceding the apostrophe.

You can use an * asterisk wildcard to begin or end the attribute value in a **Value of expression** (Equal operator only), **Check for Missing Items**, **Scan for string within a string**, or **Return a subset of the string** function.

You can also use the ? question mark wildcard for a single character at that position. When using the ? question mark wildcard with the **Value of expression** function, append the text with an * asterisk. For example, enter Sa?e* to find Safe and Sale.

Timestamp

The Timestamp attribute is treated like a text attribute formatted as mm/dd/yy hh:mm:ss. Some timestamp attributes have different names, such as *Start Date & Time* and *Recording Time*.

Some timestamp functions enable you to add and subtract values from the second attribute by using the + and - symbols. The value can be one or two digits followed by D for Day, H for Hour, M for Minute, or S for Second.

Enumerated

Enumerated attributes have a predefined set of values and the tabular editor supplies a ▾ list to choose from. These values are usually treated as text.

Some enumerated values are actually numeric, with text representing some or all of the numbers, and you can type the numeric value instead of

choosing a text equivalent from a list. For example, the WebSphere MQ attribute, *Action to Take* has values of *n/a*, *delete*, *create*, and *discover*. But you can also enter 5 or 6, which represent other actions. See the help topic of the attribute in question for a list of possible entries.

Single-row or multiple-row

Depending on the type of attribute group, data samplings return either a single row of data or multiple rows. For example, the Windows Memory attribute group returns one row of a data for a managed system, whereas the Windows Process attribute group returns multiple rows, one for each process running.

A single-row attribute group will return multiple rows if it has been configured for historical data collection.

Attributes from different attribute groups can be combined in the same situation if they are from single-row attribute groups. Local Time and Universal Time are the exception—they can be combined with any group.

When you create a query for an attribute group that returns multiple rows, consider adding a formula as a pre-filter to limit the rows returned to those of interest, which also reduces network traffic overhead.

Pure or sampled

Attributes behave differently in situations depending on whether they are able to be sampled at intervals. Attributes that cannot be sampled are those that have no value until something happens, such as those for the Windows OS File Change attribute group. These unsolicited notifications are called *pure events* Examples of pure events are an out-of-paper condition on a printer and an occurrence that adds an entry to the Windows System Log.

Events from sampled situations are closed automatically when the situation is no longer true (goes back to false). Because of the nature of pure events, they are not closed automatically like sampled events but must be closed manually unless an Until clause is included in the situation definition.

Boolean AND and OR

The tabular editors work with Boolean AND and OR logic.

Multiple expressions

Enter multiple expressions in the same row if they must all be met (Boolean AND logic) and on separate rows if any of them can be met (Boolean OR logic) to set the situation to true. For example, if you want a situation to fire when either disk time OR disk space is at least 90% on myserver, the condition would look like the one shown here.

	% Disk Time	**% Used**	**Server Name**
1	>= 90		== myserver
2		>= 90	== myserver

Branching OR expressions

If you are building a situation or setting a filter or threshold where you want an OR expression to branch at a specific AND expression, repeat the AND expression in a new row with the OR expression.

In the situation example below, the comparisons are true if the notepad process exceeds 50 seconds AND the virtual bytes exceed one million OR the notepad process on myserver exceeds 50 seconds AND the page faults exceed 100,000 per second.

	Process Name	Elapsed Time	Virtual Bytes	Page Faults/sec
1	== notepad	> 500000	> 1000000	
2	== notepad	> 500000		> 100000

Repeating attributes in AND expressions

Some situation formulas require you to repeat the same attribute in an AND expression for a certain outcome.

If you are constructing a situation that requires repeating an attribute, each instance of the attribute must appear in its own column. To repeat an attribute in a new column of the Situation editor:

- Click **Add Conditions** and select it from the Select condition window.

Examples:

- This situation is true when the notepad process occurs more than four times in the same data sampling. The user wrote the first expression to check for notepad, then clicked **Add Conditions** and selected Process Name again to create a new column and add another expression to count these processes.

	Process Name	Process Name
1	== notepad	🔢 > 4

- This situation is true when the day of the week is 02 to 06 (Monday to Friday).

	Day of Week	Day of Week
1	>= 2	<= 6

Formula editor menu

The formula editor pop-up menu has functions for editing the expressions, inserting rows, and deleting rows and columns.

Right-click a column heading, row number, or an unselected cell to see and select from a menu of actions.

Note: Although you can enter both ASCII and non-ASCII characters (such as those in double-byte character set languages) in the formula editor, you can paste only ASCII characters. If you want to repeat non-ASCII characters in a cell, you must type the characters in the cell.

Action	Description	⤻ Filters ⊟ Thresholds	🗌 Query	ƒx Formula
🖊 Clear contents	Clear the contents of the...	cell, column, or row	cell, column, or row	cell, column, or row

Action	Description	⚑ Filters ⚑ Thresholds	⬚ Query	*fx* Formula
⬚ Copy	Copy to the clipboard the contents of the...	cell or row	cell or row	cell or row
✄ Cut	Move to the clipboard the contents of the..	cell or row	cell or row	cell or row
✖ Delete	Remove the entire...	row	row	row
⬚ Edit	Open the embedded situation for editing the...			cell
⬚ Insert	Insert a new row above this...	row	row	row
⬚ Paste	Paste the clipboard contents into the...	cell or row	cell or row	cell or row

Cell functions

Use the cell formula functions to add filters to queries and query-based views, and to set thresholds for situations and query-based views. The functions that are available depend on the type of attribute (numeric, text, or time) and in which editor you are working: Query, Situation, or Properties (Filters and Thresholds tabs).

Numeric attributes
⬚ Value of expression

⬚ See if value is within a range

△ Change in value

% Percent change in value
⬚ See if Null (no value) has been detected

Text attributes
⬚ Value of expression

⬚ Return a subset of the string

⬚ Scan for string within a string

⬚ Check for Missing Items

⬚ See if Null (no value) has been detected

⬚ Scan for a string based on Regular Expression pattern

Time attributes
⬚ Return a subset of the string

⬚ Scan for string within a string

⬚ Compare date/time

⬚ Compare to time +/- delta

⬚ See if a time is within a range

⬚ Scan for a string based on Regular Expression pattern

BETWEEN: See if value is within a range

Use the BETWEEN function in the Properties editor and Query editor to filter or add a threshold indicator when a value falls within a range.

⬚ **See if value is within a range**

> Compares all values between two numbers. When you select the function, the Range Parameters window opens. Enter the first and last numbers in the value range. The relational operators available are == Equal and != Not Equal.

> ☐ **Exclude range start/end from evaluation**, when left unchecked, includes the values entered in the range. The expression ends with 0 if this is left unchecked; 1 if it is checked (see example).

Example

> The comparison is true if Copy Read Hits is between 50% and 60%.
>
> `BETWEEN (Copy Read Hits %) == 50,60,1`

Hexadecimal numbers

> When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x123ab for hex value 123ab.

CHANGE: Change in value

Use the CHANGE formula function in a situation to compare the change between the value of the current data sample and the previous sample.

△ **Change in value**

> The CHANGE function takes the amount of change between the sampled value of the attribute and the previous value, then compares it with the value you enter here. If the result meets the comparison, the situation is true. Valid values are numeric integers.

Example

> In the following example, the situation becomes true if the Virtual Bytes have increased by at least 100 bytes since the last data sampling. During the first monitor iteration the value of Virtual Bytes is 200 and during the second iteration the value is 500, which is an increase of 300 between samples, thus the situation is true.
>
> `CHANGE (Virtual Bytes) >= 100`

Restrictions

> You can have only one CHANGE function in a situation.
>
> The CHANGE function must be written on row 1 of the formula editor; all subsequent rows are disabled.
>
> You cannot combine the CHANGE function with a PCTCHANGE, MISSING or group function in a situation.
>
> A formula with the CHANGE function can use values from one attribute group only.
>
> The function list shows CHANGE for monitoring agents that are of the correct version; older versions will not display CHANGE in the list.
>
> If the expression that is responsible for a situation event uses this function, the event results workspace will show neither the attribute value nor the expression that should display when you hover over an initial situation value.

Hexadecimal numbers

> When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x123ab for hex value 123ab.

DATE: Compare Date and Time

Use the DATE formula function in a situation, query, query-based view filter, or table view threshold formula to compare the time stamp from the data sample with the value specified.

▦ **Compare date/time**

> This is the default function for time attributes. It compares the Timestamp value from the data sampling with the value entered here.

After you select the DATE function and click inside the value cell, the
Select the Time Comparison dialog opens for selecting the exact date and
time to compare. The format is the date as seen in the date columns
(mm/dd/yy hh:mm:ss) or as defined for your locale.

Example

This formula checks the Timestamp attribute value. When it matches the
date set here, the comparison is true.

```
DATE(Timestamp)
== 09/30/03 23:59:00
```

ISNULL: See if Null has been detected

Use the ISNULL formula function in a query-based view filter or a table view
threshold to find cells that have receive no value and insert null.

⊘ **See if Null (no value) has been detected**

Inserts null into a cell that has received no value from the numeric or text
attribute for the view. In the Formula editor, the value shows as True and
the operators available are == Equal and != Not Equal.

Example

In the Filters tab, any row that has no value in the Contents field will be
excluded from the view. If, instead, this formula is written in the
Thresholds tab, all cells that have a valid value in this column will be
displayed with the threshold color or threshold icon.

```
ISNULL(Contents) != True
```

MISSING: Check for missing items

Use the MISSING function in a situation to compare the data sample with a list of
text entries. If no match is found, a situation event is opened.

🔍 **Check for Missing Items**

Compares the value of the specified text attribute with a list of values you
supply. The condition is true when the value does not match any in the
list. This function is useful when you want notification that something is
not present in your system.

Examples

An event opens when the anti-virus scanning program is not available.

```
MISSING (ProcessName) == (VirusGone.exe)
```

This example builds on the last one. Two names were added to the *Missing
List* and a system name specified with a wildcard. An event will open if
any of these processes is missing on any system whose name starts with
SERVER (such as SERVER_BLDG100 or SERVERSOUTH). If you wanted to
be notified when this is the case, you could add a notification action to be
executed when the situation becomes true.

```
(MISSING (ProcessName)== (VirusGone.exe,DBServer.exe,iexplore.exe))AND
ServerName == SERVER*
```

Restrictions

You can have only one MISSING function in a situation.

The MISSING function must be written on row 1 of the formula editor; all
subsequent rows are disabled.

You cannot combine the MISSING function with any grouping functions
like MAX or COUNT but you can combine it with value-type functions
like VALUE, SCAN, DATE, etc.

Wildcards

You can use an * asterisk wildcard to begin or end the attribute value. You can also use the ? question mark wildcard for a single character at that position.

Tip   If the formula is for a multiple-row attribute group, add a display item (in the Condition tab, click **Advanced**) if you want the situation to continue testing for missing items after opening an event for the first one.

PCTCHANGE: Percent change in value

Use the PCTCHANGE formula function in a situation to compare the percent change between the value of the current data sample and the previous sample.

% Percent change in value

The PCTCHANGE function takes the percentage change between the sampled value of the attribute and the previous value, then compares it with what you enter here. If the result meets the comparison, the situation is true. Valid values are numeric integers.

Example

This situation is true when the memory usage of a process increases by 50% or more between sampling intervals.

```
PCTCHANGE (% Processor Time) >= 50
```

Restrictions

You can have only one PCTCHANGE function in a situation.

The PCTCHANGE function must be written on row 1 of the formula editor; all subsequent rows are disabled.

You cannot combine the PCTCHANGE function with a CHANGE, MISSING or group function in a situation.

A formula with the PCTCHANGE function can use values from one attribute group only.

The function list shows PCTCHANGE for monitoring agents that are of the correct version; older versions do not display PCTCHANGE in the list.

If the expression that is responsible for a situation event uses this function, the event results workspace shows neither the attribute value threshold color nor the expression that should display when you hover over an initial situation value.

Hexadecimal numbers

When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x123ab for hex value 123ab.

REGEX: Scan for a string based on Regular Expression pattern

Use the REGEX formula function in a query-based view filter or a table view threshold for the text or timestamp that matches a pattern. The pattern is based on a regular expression.

⚙ Scan for a string based on Regular Expression pattern

The REGEX function looks for a pattern match to the expression. The operators available are specific to the *regular expression* syntax. Search for "regex" or "regular expression" in your browser to find links to more information (such as this topic in the developerWorks® technical library).

Text example

> This formula for the Windows event log table view threshold highlights a cell if the event type is at least three letters long but no longer than nine. Thus, the event types *Error* and *Warning* are highlighted but not *Information*.
>
> ```
> REGEX(Type)
> == '[A-Za-z]{3,9}'
> ```

Time example

> This formula begins looking at the seventh character, the year, to find 10, so a timestamp of *06/12/10 09:29:05* would meet the criteria, but *01/20/09 11:07:05* would not.
>
> ```
> REGEX(Timestamp)
> == 7,10
> ```

Operators

> The regular expression capability in the Filters and Thresholds tabs of the Properties editor enables you to perform very specific searches through many rows of data based on a pattern. This table has a list of commonly used operators in regular expressions.

Operator	Description
. period	Match any single character
^ caret	Match the empty string that occurs at the beginning of a line or string
$ dollar sign	Match the empty string that occurs at the end of a line
A	Match an uppercase letter *A*
a	Match a lowercase letter *a*
\d	Match any single digit
\D	Match any single nondigit character
\w	Match any single alphanumeric character; a synonym is [:alnum:]
[A-E]	Match uppercase *A*, *B*, *C*, *D*, or *E*
[^A-E]	Match any characters except uppercase *A*, *B*, *C*, *D*, or *E*
X?	Match none or one capital letter X
X*	Match any number of capital *X*es
X+	Match one or more capital *X*es
X{n}	Match exactly *n* capital *X*es
X{n,m}	Match at least *n* and no more than m capital *X*es; if you omit *m*, the expression tries to match at least *n* *X*es
(abc|def)+	Match a sequence of at least one *abc* and *def*; *abc* and *def* would match

SCAN: Scan for string within a string

Use the SCAN formula function in a situation, query, view filter or table view threshold to scan the data sample for the specified text or time stamp.

abc Scan for string within a string

> Searches for the value or partial value occurring anywhere in the text or

time stamp returned by the attribute. The == Equal relational operator is valid; all other operators return invalid or unwanted results.

Text examples

This formula scans the value of Owner for *CD* anywhere in the text. For example, an owner named *ABCDEF* would meet the comparison.

```
SCAN(Owner) == CD
```

This formula scans the value of ProcessName for *ho*, so *svchost* and *shopper* would meet the comparison.

```
SCAN(ProcessName) == ho
```

Time example

This formula scans the timestamp for a value of 07/08. A date of 07/08/10 would meet the comparison as would 12/07/08.

```
SCAN(Timestamp) == 07/08
```

Wildcards

You can use an * asterisk wildcard to begin or end the attribute value. You can also use the ? question mark wildcard for a single character at that position.

STR: Return a subset of the string

Use the STR formula function in a situation, query, view filter or table view threshold to start searching for the value in the text or time stamp at a specific character position.

Return a subset of the string

Searches for a value occurring at a specific location in the data returned by the attribute. When you select STR, a Substring Comparison Information window opens for you to enter the starting position, such as 4 for the fourth character from the left to begin looking for the value.

Text example

This formula looks for ar starting with the second character in returned values for Handle of the User, so *aardvark*, *parrot* and *sardine* meet the comparison, but not *artist*, *star* or SARDINE (case-sensitive).

```
STR(Handle of the User) == 2,ar
```

Time example

This formula begins looking at the seventh character, the year, to find 10, so a timestamp of *07/12/10 09:29:05* would meet the criteria but *10/11/09 11:07:05* would not.

```
STR(Timestamp) == 7,10
```

Wildcards

You can use an * asterisk wildcard to begin or end the attribute value. You can also use the ? question mark wildcard for a single character at that position.

TIME: Compare to a time +/- delta

Use the TIME formula function in a situation, query-based view filter, or table view threshold formula to compare the time stamp attribute (such a recording time) with the specified difference in number of days, hours, minutes, or seconds.

Compare to time +/- delta

Compares the time difference you enter with the timestamp returned by an attribute. After you select the TIME function and click inside the value cell,

the Select Time Comparison Criteria window opens. Select a timestamp attribute from the list. Specify the delta by combining + or - with the number of days (D), hours (H), minutes (M), or seconds (S). The sampled timestamps are compared to the delta.

Example

This formula compares the time the event was logged (Entry Time) with the timestamp from the data sampling. If the event occurred seven days earlier, the comparison is true. If the == were changed to <=, the comparison would be true after eight days, nine days, and so on.

```
TIME(Entry Time) == #NT_Event_Log.Timestamp - 7D
```

TIMESPAN: See if a time is within a range

Use the TIMESPAN formula function to filter a query-based view or to add threshold indicators to a table view. The data sample time stamp is compared with the specified time span.

🕐 **See if a time is within a range**

The TIMESPAN function selects rows based on the comparison of the Timestamp value to the time span indicated. When you select this function and click inside the value cell, the **Select Time Span** dialog opens for you to choose a time span. Note that only the == and != relational operators are available for this function.

Example

This formula selects all rows with timestamps within the last two hours.

```
TIMESPAN(Timestamp)
== ,,7200
```

VALUE: Value of expression

Use the VALUE formula function in a situation, query, view filter or table view threshold to look for the partial or complete number or text in the data sample.

▼ **Value of expression**

Takes the sampled value of the attribute. When used with other functions, VALUE is evaluated before the others, acting as a filter to the other functions.

Numeric example

This formula looks through all CPU Utilization values. If one is over 50% the comparison is true.

```
VALUE(CPU Utilization)> 50
```

Text examples

The comparison is true if the system name is *redwood*.

```
VALUE(System Name) == redwood
```

This formula uses a wildcard to find log files that start with uppercase *S*, such as *System* and *Security*.

```
VALUE(Log Name) == S*
```

Wildcards

For text attributes only, you can use wildcards to enter partial values that will find everything that includes those characters provided that you append the text with an * asterisk. Use an * asterisk wildcard for any number of characters beginning at that position. Use a ? question mark wildcard for a single character at that position. The last character of the text must be an * asterisk. For example, enter Sa?e* to find Safe and Sale;

and enter *System* to find AIX_System and UNIX_Systems. Be aware that only the == Equal operator is valid; all other operators will return invalid or unwanted results. For another example, db2*.exe* will find db2fmp.exe and db2systray.exe; and db2???.* will find db2fmp.exe but not db2systray.exe.

Hexadecimal numbers

When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x123ab for hex value 123ab.

Tip For situations on multiple-row attribute groups, you can add a display item (in the Condition tab, click **Advanced**) to open an event for each row that meets the condition and not just the first one.

IN: Value of expression IN

Use the IN relational operator with the 𝑽 Value of expression function in a situation or a query-based view filter or threshold to compare the data sample with a list of text values. The condition is met when a match is found.

𝑽 Value of expression IN

Compares the value of the specified text attribute with a list of values that you supply. The condition is true when a match is found in the list. Using the IN relational operator with the 𝑽 **Value of expression** function is useful for highlighting specific names or text values that are present in the data samplings, such as certain process names.

Examples

This is a filter used to display only those rows that have a user name of Anne, Jane or John in the table.

```
(User Name *IN ('Anne','John','Jane'))
```

This is a threshold formula for a view of situations used to apply the Minor severity threshold color to any cell that contains one of these names.

```
((--- Minor ---) Name *IN ('NT_Percent_Processor_Time',
  'NT_Memory_Pages_Sec_Low','NT_Percent_Disk_Time_Low'))
```

Situation tip

If the formula is for a multiple-row attribute group, add a display item (in the Condition tab, click **Advanced**) if you want the situation to continue testing for missing items after opening an event for the first one.

Group functions

Use the group functions in the Query editor or Situation editor for attribute groups that return multiple rows of data or a single-row attribute group that has been configured for historical data collection. With the exception of COUNT, which can be used by all attribute types, group functions are available only for numeric attributes.

Note:

1. ▣ Time Span selection is not available for a table or chart view if it uses a query that includes a group function.
2. When entering a hexadecimal value in a cell, prefix the number with zero and the letter "x", as in 0x456c for hex value 456c.
3. If the expression that is responsible for a situation event uses a group function, the situation event results workspace shows neither the attribute value threshold color nor the expression that should display when you hover over an initial situation value.

AVG: Average of group

Use the ▤ **Average of group** function in a query or situation to calculate the average from a range of values.

▤ **Average of group**

The AVG function totals the attribute values specified and calculates the average.

Query example

This function averages all returned values of % User Time. GROUP BY is Server Name, so the view will show the average user time, one row for each managed system.

```
AVG(% User Time)
```

Situation example

This formula averages all returned values of % User Time. If the average exceeds 20%, the comparison is true.

```
AVG (% User Time)> 20
```

COUNT: Count of group members

Use the ▦ **Count of group members** formula function in a query or situation to count the number of returned rows, which can be filtered by an AND expression.

COUNT in a query to the monitoring server works differently than it does in a situation: It counts the column specified within the group referenced (Advanced button, Group By tab), and displays the result in the count column.

Note: When using COUNT in an AND expression, it must be the last function in the expression.

Query example

This function counts the number of processes. GROUP BY is Server Name, and the number of processes counted for each system appears in the ID Process column. Example as it would be written in an SQL statement:

```
COUNT(ID_Process)
```

Situation examples

Examples as they appear in ⚙ Show Formula:

- The comparison is true when more than 10 rows have been counted.
  ```
  COUNT(Timestamp) > 10
  ```

- This formula first filters the rows for those with the notepad process, then counts the returned rows by referencing the Timestamp attribute. The attribute used for the count does not matter, because the COUNT function is simply counting the number of rows where this attribute appears. When the count exceeds four, the comparison is true.
  ```
  (Process Name) == notepad AND COUNT(Timestamp) > 4
  ```

- The next formula counts the processes that have been running for more than an hour. When at least five processes have been counted, the comparison is true.
  ```
  (Elapsed Time (Seconds) >= 3600)) AND COUNT(Timestamp) >= 5
  ```

MAX: Maximum in group

Use the ▤ **Maximum in group** function in a query or situation to compare the attributes returned and select the one with the highest value.

■ **Maximum in group**

> For attributes that return multiple rows, the MAX function finds the highest returned value (or values).

Query example

> This function finds the process with the highest thread count. GROUP BY is Server Name, so the view shows the process with the highest thread count on each system.
>
> `MAX(Thread Count)`

Situation example

> The relational operator is set to == and the value to TRUE. The MAX function is meaningful when combined with another expression. This formula finds the process with the highest thread count. If the process also has more than 150 million virtual bytes, the comparison is true.
>
> `MAX(Thread Count)== *TRUE AND Virtual Bytes >= 150,000,000`

MIN: Minimum in group

Use the ■ **Minimum in group** function in a query or situation to compare the attributes returned and select the one with the lowest value.

■ **Minimum in group**

> For attributes that return multiple rows, the MIN function finds the lowest returned value (or values).

Query example

> This function finds the process with the lowest priority base. GROUP BY is Server Name, so the view will show the process with the lowest priority base on each system.
>
> `MIN(Priority Base)`

Situation example

> The relational operator is set to == and the value to TRUE. The MIN function is meaningful when combined with another expression. This formula finds the process with the lowest priority base. If it also has more than 20% privileged time, the comparison is true.
>
> `MIN(Priority Base)== *TRUE AND % Privileged Time >= 20`

SUM: Sum of group

Use the Σ **Sum of group** function in a query or situation to calculate the total from a range of values.

Σ **Sum of group**

> Takes the attribute values returned that meet the criteria and totals them.

Query example

> This formula totals all CPU Utilization values in the data sampling. GROUP BY is Server Name, so the view will show the total for each managed system.
>
> `SUM(CPU Utilization)`

Situation example

> This formula totals all CPU Utilization values in the data sampling from the system named *redwood*. If the sum of these values exceeds 50%, the comparison is true.
>
> `SUM(CPU Utilization)> 50 AND System Name == redwood`

Appendix B. Tivoli Enterprise Monitoring Server and global attributes

The Tivoli Enterprise Monitoring Server has attributes for monitoring enterprise logs and status. Global attributes are also available for use alone or combined with other attribute groups in situations.

Every monitoring agent connects to a Tivoli Enterprise Monitoring Server, which acts as a collection and control point for alerts received from the agents, and gathers performance and availability data from them. The main, or hub, monitoring server correlates the monitoring data collected by agents and remote servers and passes it to the Tivoli Enterprise Portal Server for presentation in the Tivoli Enterprise Portal.

Monitoring agents are made up of attributes, which represent the properties of systems or networks, such as the amount of CPU usage or the message ID. Attributes are organized into attribute groups. The attributes in a group can be displayed in a query-based view (table, chart or relational topology) or used to specify a condition for testing in a situation. When you open the view or start the situation, data samples are taken of the selected attributes.

The Tivoli Enterprise Monitoring Server has attribute groups that you can use for reporting the status of enterprise activities, monitoring agents, and situations:

- Agent Operations Log for use in query-based views and situations to monitor message-related information about the monitoring agent. See the "Agent Operations Log attributes" on page 432.
- Three attribute groups for the Depot
- Enterprise Information Base (EIB) Change log for use in query-based views of the changes made by users to deploy situations and policies. See the "EIB Change Log attributes" on page 440.
- Managed System for use in query-based views and situations regarding the status of the monitoring server and agents. See the "Managed Systems attributes" on page 449.
- Situation Definition for use in query-based views to report the complete definition of situations. See the "Situation Definition attributes" on page 450.
- Active Situation Status and Situation Status for use in query-based views to report the situation and policy status. See the "Situation Status Current attributes" on page 453 and "Situation Status Log attributes" on page 455.
- Application Property Installation provide information related to the self-describing agent installation process. See the "Application Property Installation attributes" on page 432.

The Tivoli Enterprise Monitoring Server also has several global attribute groups for using alone or combining with a product attribute group to create situations:

- Local Time and Universal Time for use in situations, including embedded situations on any managed system. See the "Local Time attributes" on page 448 and "Universal Time attributes" on page 457.
- Universal Message for monitoring the messages that have been generated through a situation or policy action command. See the "Universal Messages attributes" on page 456.

- The Situation Status Event attributes for information on situation events. Views that display this information include My Acknowledged Events, Selected Event, and Similar Events and are used in some of the Enterprise workspaces.
- ITM Audit attributes for information on critical state-changing events in the monitoring environment. See the "ITM Audit attributes" on page 441.
- ITM Historical Collection attributes for information about the active historical collections that are exporting data from a location. See the "ITM Historical Collection attributes" on page 446.

You can also work with attributes of other applications by creating SQL queries to retrieve them from ODBC-compliant databases.

Agent Operations Log attributes

The Agent Operations Log collects the messages that occur on managed systems in your enterprise. This log provides a list of message-related information about the monitoring agent.

Restriction: The IBM i and z/OS monitoring agents do not list information in the Agent Operations Log. Also, because this log is specific to the individual agent, it reports data only at the ⚙ agent and 🖥 attribute levels of the Navigator Physical view.

Global Timestamp The date and time at the monitoring server when the operation took place. The format is MM/DD/YY HH:MM:SS.

Help Text The text description of the operation.

Managed System Type The type of managed system.

Message Number The message ID that corresponds to the text.

Message Severity A numeric indicator of the severity of the operation.

Message Text The operation that took place at the agent, such as execution of a take action command or connection to the monitoring server.

Server Name The name of the monitoring server where the operation took place.

Application Property Installation attributes

Application Property Installation attributes provide information related to the self-describing agent installation process. This table stores and maintains: the status of each product's self-describing agent product support installation, the operational state of the self-describing facility for each monitoring server, and user configuration settings for the self-describing feature. This attribute group is also known as the TAPPLPROPS table.

Note: Some attributes might return blank data because they are not used by a specific record type or because they are reserved for future use.

Restriction: This attribute group is only available for use in queries.

Catalog Install Timestamp The package record date and time that is extracted from the product catalog file (kxx.cat) on the monitoring server. The format is CYYMMDDHHMMSSmmm (as in 1120615074501000 for June 15, 2012 at 07:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m = Millisecond

Catalog Install Version The package record version that is extracted from the product catalog file (kxx.cat) on the monitoring server. This version is not expected to exactly match any self-describing reported product version.

Catalog SDA The value identifies the results of the self-describing consistency check performed at the monitoring server startup. Determines whether a manual install process was performed after the last monitoring server self-describing installation. Valid values:

Value	Description
Y	Self-describing check passed. The stored values for the **Catalog Install Timestamp** and **Catalog Install Version** attributes match the values from the existing product catalog files that are installed on the monitoring server.
N	Self-describing check failed. A mismatch was detected between the **Catalog Install Timestamp** and **Catalog Install Version** attributes, as compared to the existing product catalog files that are installed on the monitoring server, due to an undetermined manual install process.

Configuration The configuration details related to the type of record being viewed.

Group ID The value identifies the record type. The value 5655 identifies a self-describing product installation record. The value 5530 identifies a self-describing configuration record.

ID The value indicates the type of record.

The following values are the current package keys for self-describing installation records:

Key	Description
TMS	Monitoring server product support package
TPS	Portal server product support package
TPW	Portal browser client product support package

The following values are the current package keys for self-describing user or status configuration types:

Key	Description
SEEDDIST	Product seeding configuration records
SDA_STAT	Monitoring server self-describing operation status records
INSTALL	Product install user configuration records
SUSPEND	User suspended operation records

ID Version The version of the ID column key for the self-describing product installation status records. The format is: VVRRMMFF (as in 06230100 for Version 6.2.3 Fix Pack 1) where:

> V = Version
>
> R = Release
>
> M = Modification
>
> F = Interim Fix

Last Date The date and time of the last recorded update. The format is CYYMMDDHHMMSSmmm (as in 1120615074501000 for June 15, 2012 at 07:45:01) where:

> C = Century (0 for 20th, 1 for 21st)
>
> Y = Year
>
> M = Month
>
> D = Day
>
> H = Hour
>
> M = Minute
>
> S = Second
>
> m = Millisecond

Manual Install The value indicates a manual product install for self-describing product installation status records only. Valid values:

Value	Description
Y	The monitoring server support was applied by a manual installer.
N	The monitoring server support was installed by the self-describing process.

Product The product code as reported by the self-describing product manifest file (such as `NT` or `LZ`). Typically returns the two-character product code. Some configuration records types return a value of `TEMS`.

Product Version The version of the **Product** attribute for self-describing installation status records. The format is: VVRRMMFF (as in 06230100 for Version 6.2.3 Fix Pack 1) where:

> V = Version
>
> R = Release
>
> M = Modification
>
> F = Interim Fix

Resource The self-describing package file name used for the product installation.

Seeding State The product seed processing state or result of any product seed operation that is invoked directly by the monitoring server self-describing agent installation process. Valid values:

Value	Description
Y	Seed operation completed successfully.
N	Seeding has not run.
E	An error occurred during the seeding operation.
I	Seed operation started.
U	State of seed operation is unknown and displays as blank. This is the initial value.

State The processing state codes for the self-describing product installation status. Valid values:

Value	Description
IR	Install request
IM	Install metadata
MC	Metadata complete
ME	Metadata error
IS	Install seed data
SC	Seed data install complete
SE	Seed data error or incomplete
IC	Installation complete
	(Blank) Product installed manually and not by self-description.

Status The operation or phase completion status, if any. A zero (0) indicates that self-describing is enabled.

Timestamp The date and time the TAPPLPROPS record was created. The format is CYYMMDDHHMMSSmmm (as in 1120615074501000 for June 15, 2012 at 07:45:01) where:

C = Century (0 for 20th, 1 for 21st)

Y = Year

M = Month

D = Day

H = Hour

M = Minute

S = Second

m= Millisecond

Catalog Check attributes

The Catalog Check attributes compare the application catalogs of the hub Tivoli Enterprise Monitoring Server with remote monitoring servers and report any discrepancies. These are the predefined queries for Catalog Check attributes: *Compare Enterprise Catalogs* and *Compare Tivoli Enterprise Monitoring Server Catalogs.*

Application Name The identifier of an application catalog that is installed at the Tivoli Enterprise Monitoring Server. Some of the catalog IDs are identified in the *IBM Tivoli Monitoring Messages Guide*.

Hub Column Count The total number of table columns for this application at the Hub Tivoli Enterprise Monitoring Server. The catalog contains all the agent table definition, with multiple tables of multiple columns.

Hub Timestamp The timestamp of the application catalog at the Hub Tivoli Enterprise Monitoring Server. The format is MM/DD/YY HH:MM:SS.

Server Column Count The total number of table columns for this application. If the total is different from the hub column count, it might signify that the version installed is different from the version of application support installed on the hub monitoring server.

Server Name The name of the Tivoli Enterprise Monitoring Server.

Server Timestamp The timestamp of the application catalog at the Tivoli Enterprise Monitoring Server. The format is MM/DD/YY HH:MM:SS.

Status The result of the comparison between an application catalog at the Hub Tivoli Enterprise Monitoring Server and the corresponding application catalog at a Remote Tivoli Enterprise Monitoring Server.

Enumeration	Description
0	**OK** indicates that the timestamps and column counts of the hub and remote monitoring servers match.
1	**Hub Newer** indicates that the hub monitoring server was updated more recently than a remote monitoring server.
2	**RTEMS Newer** indicates that a remote monitoring server was updated more recently than the hub monitoring server.
3	**Missing At Hub** indicates that the catalog ID that is present on a remote monitoring server is missing at the hub monitoring server.
4	**Missing At RTEMS** indicates that the catalog ID that is present at the hub monitoring server is missing at a remote monitoring server.
5	**Count Mismatch** indicates that the column count at the remote monitoring server does not match the column count at the hub monitoring server.
6	**Timestamp Mismatch** indicates that the time the remote monitoring server was last updated does not match the time the hub was last updated.

Timestamp The time at the Tivoli Enterprise Monitoring Server location when the data was sampled. The format is MM/DD/YY HH:MM:SS.

Deploy Status attributes

The Deploy Status attributes give information on the status of the agent deployment; each row displays a completed command.

Bundle group The collection of agents, another bundle group, or collection of bundle groups that are being added to the computer.

Command The command being used for the agent deployment process, such as updateAgent.

Command	Description
addSystem	The command to deploy a monitoring agent to a computer in the IBM Tivoli Monitoring environment. This command deploys an agent and other required components if they are not already installed on the node.
createNode	The command to create the node directory on the target computer. This is the directory where the OS agent is installed and where any non-OS agents are deployed.
setConfig	The command to edit configuration options of an existing agent. The agent is restarted so that the new configuration parameters can take effect. Specify the configuration data through the parameter pair *section.name=value*.
startAgent	The command to start the monitoring agent.
stopAgent	The command to stop the monitoring agent.
updateAgent	The command used to install an agent update on a specified managed system. Updating agents involves stopping any that are running, applying the changes, and restarting them. If a version is not specified, the agent is updated to the latest version.

Count Counts to keep track of number of requests from the group.

Deploy group The collection of host names of different computers or other deploy groups that the agents are deployed to. The deploy group will have a deploy name.

Deploy group name The name of the group containing the collection of host names of different computers or other deploy groups that the agents are deployed to.

Deploy TEMS The through node or the remote Tivoli Enterprise Monitoring Server that the agent is being deployed from.

GlobalTimeStamp The most recent update at the hub Tivoli Enterprise Monitoring Server to the deployment status for the transaction time. The GlobalTimeStamp attribute indicates the last time any information was updated for the deploy command, which is used for EIB probe completion.

Group Transaction ID The identifying number that is assigned to a group transaction. This value allows you to see the transactions involved in the same command. It is the amef value.

lstdate timestamp The last time the record for the deploy command (row in the table) was changed. The lstdate timestamp attribute displays the last time any information update made. The files have been installed and the agent is configured; the row in the table is inserted once the command is complete.

NodeId The agent name and platform (separated by a colon). The nodeId and the target managed system name are similar.

Platform Version The operating platform of the target computer for the agent install. For example, WinNt.

Processing State The state of the deploy command, from pre-deployment to post-installation. The state is indicated with an integer.

Product ID The product code of the agent.

Product Version The version of the agent.

Reserved1 Reserved for future use.

Reserved2 Reserved for future use.

RetryCount The number of times the agent is retried. 0 is displayed if it is fine and no retries occur.

Return code The code that is generated upon completion of the task: 0 for success; 3 for failure.

Return Message The information message or error message regarding the agent deployment. For example, fail is displayed when agent has been deployed and couldn't connect.

StartTime The time of the most recent attempt to initiate the deployment transaction. With every retry, this value is reset. No value is displayed if no retry occurs.

Status The status of the agent deployment. Valid values are Success, Failed, In progress, and Failed retry (default setting for number of attempts to retry is 5 times with 90 second intervals between each retry).

System The host address of the deploy target computer. It is the fully qualified host name or IP address of the target computer.

Target managed system name The managed system name of the OS agent that is processing the deployment transaction.

Transaction ID The identifying number assigned to the transaction; this identifier is unique to each transaction. It is generated through a random string of numbers.

Deploy Summary attributes

The deploy summary attributes provide a summary of all the deploy commands that have been issued.

Bundle group The collection of agents, another bundle group, or collection of bundle groups that are being added to the computer.

Command The commands that are used for agent deployment, maintenance, and system configuration.

Command	Description
addSystem	The command to deploy a monitoring agent to a computer in the IBM Tivoli Monitoring environment. This command deploys an agent and other required components if they are not already installed on the node.
createNode	The command to create the node directory on the target computer. This is the directory where the OS agent is installed and where any non-OS agents are deployed.
setConfig	The command to edit configuration options of an existing agent. The agent is restarted so that the new configuration parameters can take effect. Specify the configuration data through the parameter pair *section.name=value*.
startAgent	The command to start the monitoring agent.
stopAgent	The command to stop the monitoring agent.
updateAgent	The command used to install an agent update on a specified managed system. Updating agents involves stopping any that are running, applying the changes, and restarting them. If a version is not specified, the agent is updated to the latest version.

Deploy group The collection of host names of different computers or other deploy groups that the agents are deployed to. The deploy group will have a deploy name.

Deploy TEMS The through node or the remote Tivoli Enterprise Monitoring Server that the agent is being deployed from.

Failed The number of deploy commands that have failed.

Group Transaction ID The identifying number that is assigned to a group transaction. This value allows you to see the transactions involved in the same command.

In Progress The number of deploy commands that are in progress.

Product ID The agent product code.

Queued The number of deploy commands that are queued.

Retrying The number of deploy commands that are retrying.

Succeeded The number of deploy commands that have succeeded.

Summarize On The selection that distinguishes how the table is ordered.

Depot Inventory attributes

The depot inventory attributes give information on agent bundles in the depot.

Config descriptor version The version of the configuration descriptor.

Description The description of the agent bundle.

Friendly Version The more user friendly version of the agent bundle.

Host type The operating platform of the bundle can be installed on.

Host versions The specific version associated with the operating platform the bundle can be installed on.

Prerequisites The prerequisites required for this bundle to be installed properly.

Product ID The two character code of the agent.

Refresh Depot The value describing whether the depot information needs to be refreshed.

Version The version of the agent.

EIB Change Log attributes

The Enterprise Information Change (EIB) log records the details of changes made by users to deploy situations and policies.

Global Timestamp The date and time at the monitoring server when the change took place. The format is MM/DD/YY HH:MM:SS.

Last User The user ID of the person who made the change, up to 10 letters and numbers.

Object Name The name of the situation, policy, or other entity that was changed.

Operation The type of operation that took place. This is an enumerated attribute where a single word conveys the full meaning of the operation:

Add	A managed system or managed system group was added to the situation or policy distribution.
Change	The situation was distributed and is running.
Delete	A situation or policy was deleted from the monitoring server.
Erase	A managed system or managed system group was removed from the situation or policy distribution.
Insert	A new situation or policy was added (inserted into the internal database.
Update	A situation or policy was updated.

Path Name The location of the file that was changed, up to 96 characters.

Server Name The name of the monitoring server where the change took place.

Table Name The name of the database table on the monitoring server where the change was generated. The enumerated values are the class IDs for the objects.

5535	IDDefAccessList	The distribution list for the policy or situation has changed.
5529	IDDefNodeList	Managed system group.
5130	IDDefPolicy	The policy definition.
5535	IDDefSituation	The situation definition.
5625	IDRestartCommand	The situation has been restarted.

ITM Audit attributes

The ITM Audit attributes give information on critical state-changing events in the monitoring environment. This attribute group includes information such as which user, application, or process initiated the event; what type of action the event represents; when did the event happen; what database, application, or permission was manipulated; on which computer did this event happen; which system is the source of the event; and which system is the target of the event. You can create Enterprise level queries and situations with this attribute group.

Note: Some attributes might return blank data because they are reserved for future use.

Situations created using this attribute group are displayed under All Managed Systems, *not* under the node that you created the situation from. Pure events are generated for situations created with this attribute group.

Application Code The product code for the application. Character limit: 16.

Application Component The specific component within the product that generated the event. Character limit: 16.

Application Version The version of the application. Character limit: 24.

Assumable Authorization ID The authorization ID granted to the user or entity. Character limit: 128.

Audit Record Version The version of the audit information. Character limit: 4.

Authorization ID The role, group, or user the User ID or Entity is authorized under. Character limit: 128.

Authorization Plugin Type The type of authorization plug-in in use at the time the audit event occurred. Reserved for future use. Character limit: 4.

Enumeration	Description
-1	Not_Applicable
0	LOCAL
1	SSO
2	LDAP
4	TSPM
16	RBAC

Enumeration	Description
32	DIS
128	TIM
256	WAS
512	OTHER
1024	TAM

Correlator A unique number to correlate all the audit records related to a single event. Reserved for future use. Character limit: 8.

Enumeration	Description
-1	Not_Available

Domain The multi-tenant domain identifier. A custom identifier to identify unique namespaces within the product namespace. Typically used for separating customers logically within a common infrastructure (ISP Model). This is set by providing a value for the ITM_DOMAIN environment parameter. Character limit: 128.

Entity Identifies a non-user entity that initiated the audit event. Character limit: 128.

Entity Type Describes the type of entity that initiated the Audit Event. Character limit: 128.

Extra Attributes A string that contains a semi-colon delimited list of name="value" pairs for which no dedicated column exists in the ITM Audit table. The event's message parameters are stored in the Extra Attributes field as ParamN=n;PARM1="VALUE";PARM2="VALUE"...;PARMn="VALUE". In the name="value" pairs, the value is always enclosed in quotation marks (" ") and embedded quotation marks are escaped by doubling them. The separator between name="value" pairs is a semicolon (;). No whitespace is allowed around the equal sign (=) or semicolon. **Note:** The column can hold only 512 bytes, so the string is truncated at 512 characters if longer than that. There will be fewer than 512 characters if multi-byte characters are used.

Grantee The grantee ID for which a privilege or authority was granted, revoked or checked for authorization. Reserved for future use. Character limit: 128.

Grantee Type The Common Data Model type of the grantee. Reserved for future use. Character limit: 128.

Message The audit event descriptive message text. Character limit: 512.

Object Name The descriptive name of the object. Character limit: 256.

Object Path The detailed path to the object. Character limit: 256.

Object Type The Common Data Model class of the object. Public CDM classes are defined in CDM v2.10.21 or greater. Private CDM classes are prefixed with ibm-prod-tivoli-itm: private namespace designation. Character limit: 128.

Object Version The version of the object. Character limit: 24.

Operation Object Type The Common Data Model class of the operation. Public CDM classes are defined in CDM v2.10.21 or greater. Private CDM classes are prefixed with ibm-prod-tivoli-itm: private namespace designation. Character limit: 128.

Operation Type The Common Data Model action of the operation. The action type should be compatible with the type of object involved. Public CDM Actions are defined in CDM v2.10.21 or greater. Character limit: 128.

Operation Name The descriptive name of the operation. Character limit: 128.

Origin The network name of the managed system that emitted the audit event. Character limit: 128.

Origin Address The address of the node that emitted the audit event. Character limit: 64.

Origin Hostname The hostname of the node that emitted the audit event. Character limit: 128.

Origin Name The descriptive name of the node that emitted the audit event. Character limit: 128.

Origin Port The communication port used by the origin node. Character limit: 4.

Enumeration	Description
-1	Not_Available

Origin Protocol The communication protocol used by the origin node. Character limit: 16.

Origin Type The descriptive classification of the node that emitted the audit event. Character limit: 4.

Enumeration	Description
-1	Not_Available
1	GUI
2	CLIENT
3	SERVER
4	ENDPOINT
5	AGENT
6	CLI
7	OTHER

Privilege The access type for which a security label is granted. The set of values possible here are specified as the set of actions appropriate to the type of object specified in the object parameters. Reserved for future use. Character limit: 128.

Privilege Type Indicates the type of privilege or authority granted or revoked. Reserved for future use. Character limit: 128.

Appendix B. Tivoli Enterprise Monitoring Server and global attributes **443**

Record Type Indicates the set of attributes expected in the audit event. Character limit: 4.

Enumeration	Description
32	Permission_Checking
33	Object_Maintenance
34	Security_Maintenance
35	System_Administration
36	Contextual_Event
37	Authentication_Validation

Resource Bundle Key The message ID number. Character limit: 64.

Result The application specific return code. This allows for easier filtering and situation creation. Character limit: 4.

RunAs The role, group, or user the user ID is authorized under. Character limit: 128.

Security Policy Name The identifier for the security policy authorizing this action. Reserved for future use. Character limit: 128.

Sequence Number The audit event count of this particular audit event since the component was started. The audit event count is reset to zero when the component is reset. Character limit: 8.

Enumeration	Description
-1	Not_Available

Service Point The origin's service point. Refer to "Using the IBM Tivoli Monitoring Service Console" in the IBM Tivoli Monitoring Information Center. Character limit: 128.

Source The Managed System Name of the node that initiated the operation reported by the audit event. Character limit: 128.

Source Address The address of the node that initiated the operation reported by the audit event. Character limit: 64.

Source Hostname The hostname of the node that initiated the operation reported by the audit event. Character limit: 128.

Source Name The descriptive name of the node that initiated the operation reported by the audit event. Character limit: 128.

Source Port The communication port used by the source node. Character limit: 4.

Enumeration	Description
-1	Not_Available

Source Protocol The communication protocol used by the Ssource node. Character limit: 16.

Source Type The descriptive classification of the node that initiated the operation reported by the audit event. Character limit: 4.

Enumeration	Description
-1	Not_Available
1	GUI
2	CLIENT
3	SERVER
4	ENDPOINT
5	AGENT
6	CLI
7	OTHER

System Name The Managed System Name of the node that emitted the audit event. Character limit: 32.

Target The Managed System Name of the node where the operation is targeted. Character limit: 128.

Target Address The address of the node where the operation is targeted. Character limit: 64.

Target Hostname The hostname of the node where the operation is targeted. Character limit: 128.

Target Name The descriptive name of the node where the operation is targeted. Character limit: 128.

Target Port The communication port used by the target node. Character limit: 4.

Enumeration	Description
-1	Not_Available

Target Protocol The communication protocol used by the target node. Character limit: 16.

Target Type The descriptive classification of the node where the operation is targeted. Character limit: 4.

Enumeration	Description
-1	Not_Available
1	GUI
2	CLIENT
3	SERVER
4	ENDPOINT
5	AGENT

Enumeration	Description
6	CLI
7	OTHER

Timestamp This is the time the log event was generated. The format is MM/DD/YY HH:MM:SS. Character limit: 16

Timestamp (MS) The UTC Timestamp in milliseconds from the epoch. (Number of seconds elapsed since midnight Coordinated Universal Time (UTC) of January 1, 1970, not counting leap seconds. Also known as POSIX time or UNIX time.) Character limit: 8.

Trace Level The minimum AUDIT_TRACE level required to generate this message. Character limit: 4.

Enumeration	Description
4000	Detail
5000	Basic
6000	Minimum

User ID The user ID that initiated the operation. Character limit: 128.

ITM Historical Collection attributes

ITM Historical Collection attributes provide information about the active historical collections that are exporting data from a location, including the metrics about that historical collection.

Attribute Group Name The name of the attribute group for which the collection is being performed.

Collection Identifier The name of the situation that triggered the collection.

Collection Interval The amount of time between collections of historical data, in seconds. This attribute is only relevant for sampled data, not event data.

Collection Type Indicates the type of collection, enterprise or private.

Enumeration	Description
1	Enterprise
2	Private

Data Exported The total amount of uncompressed data exported, in KB.

Export Duration The amount of time that the export of historical data took, in seconds.

Export End The time that the export of historical data finished.

Export Interval The amount of time between exports of historical data, in seconds.

Export Interval Used The percentage of the time between exports of historical data used by this export.

Export Start The time that the export of historical data began running.

Export Status Indicates if the export of historical data is idle, queued, trimming data, or processing rows.

Enumeration	Description
1	Request Queued
2	Trimming Data
4	Processing Rows
4	Idle

Last Export Return Status The return status of the most recent export. This is the warehouse proxy return code.

Last Successful Export Interval The amount of elapsed time since the last successful export of historical data, in seconds.

Last Successful Export Time The end time of the last successfully completed export.

Processing Duration The amount of time that the export of historical data took to process rows, in seconds.

Processing End The time that the export of historical data finished processing rows.

Processing Start The time that the export of historical data began processing rows.

Product Code The product code for the historical collection.

Row Size The size of each row fetched, in bytes.

Rows Corrupted The number of rows that could not be processed or skipped due to errors.

Rows Processed The number of rows successfully exported to the Warehouse Proxy Agent.

Rows Skipped The number of rows skipped.

Rows Total The total number of rows fetched.

Sample Type Indicates the sample type, sampled or event.

Enumeration	Description
1	Sampled
2	Event

System Name The name of the system where the record originated from.

Table Name The name of the table for which the collection is being performed.

Trim Duration The amount of time that the export of historical data took to trim data, in seconds.

Trim End The time that the export of historical data finished trimming data.

Trim Start The time that the export of historical data began trimming data.

Local Time attributes

The Local Time attributes show the time and date where the agent is located. For example, if you are running the Tivoli Enterprise Portal in London and monitoring a system in Cairo, the local Cairo time and date will be returned. These attributes are useful for action situations triggered by time.

Note:
1. The Local Time attribute group can be combined with any attribute group to create a situation condition.
2. The Local Time attributes are not available for use in Tivoli Enterprise Monitoring Server queries.
3. Seconds, Minutes, Hours, and Day Of Month are treated as two-character text attributes. For example, *02* Hours is correct, but *2* Hours will never be true.

Day Of Month The day in the month that the Tivoli Enterprise Monitoring Server samples the data at the agent location, from 00 to 31. For example, 01 is the first of the month.

Day Of Week The day in the week that the Tivoli Enterprise Monitoring Server samples the data at the agent location, from Sunday to Saturday. Select the day from the ▾ list.

Hours The hour in the day that the Tivoli Enterprise Monitoring Server samples the data at the agent location, from 00 to 23, starting at 00 for 12 am.

Minutes The minutes into the hour that the Tivoli Enterprise Monitoring Server samples the data at the agent location, from 00 to 59, such as 05 for five minutes past the hour.

Month Of Year The month in the year that the Tivoli Enterprise Monitoring Server samples the data at the agent location. Select the month from the ▾ list.

Seconds The second in the minute that the Tivoli Enterprise Monitoring Server samples the data at the agent location. It can be from 00 to 59, such as 08 for the eighth second.

Time The time of the data sampling, corrected for local time zone and daylight saving time, formatted as HHMMSS. For example, 170700 is 5:07 PM.

Timestamp The date and time when the Tivoli Enterprise Monitoring Server samples the data at the agent location. The format is MM/DD/YY HH:MM:SS.

Managed Systems attributes

The Managed Systems attributes are used to check vital heartbeat information from the Tivoli Enterprise Monitoring Server and monitoring agents. You can apply the Managed System query to a view (in the Query editor under Tivoli Enterprise Monitoring Server → Managed System).

For the Navigator Physical view Enterprise item and for custom Navigator view items that have the *HUB_ or *ALL_CMS managed system group assigned, you can also create situations with these monitoring server attributes.

Affinities The code identifier of the product component.

Expiry Time The time when the hub Tivoli Enterprise Monitoring Server stops waiting for the heartbeat and considers the agent offline. This attribute is available for queries, but not for situations. The format is CYYMMDDHHMMSSmmm (as in 1090415074501000 for April 15, 2009 at 07:45:01) where:

    C = Century (0 for 20th, 1 for 21st)

    Y = Year

    M = Month

    D = Day

    H = Hour

    M = Minute

    S = Second

    m = millisecond

Host Address The IP address and port number of the monitoring agent. Other address types can be IP pipe and SNA.

Host Info The name and version of the operation platform and any applied Service Packs. This attribute is available for queries, but not for situations.

Managing System Also called through node, this is the parent system that collects data from the agent or subagent. If the agent is directly connected to the hub Tivoli Enterprise Monitoring Server, the hub monitoring server is the managing system; if the agent is connected to a remote monitoring server, then that is the managing system. An agent with subagents is the managing system for each subagent.

Name The host name of the agent.

Origin Node The network name of the source for this information. Valid entries are up to 32 characters. This attribute is available for queries, but not for situations.

Product The two-letter code of the monitoring agent.

Reason The two-character reason code, if any, for the status. For example, FA means that the agent generated its own offline message, and 00 means that the heartbeat is successful. Reasons are related to the communications layer.

Status Indicates whether the monitoring agent is ONLINE (Y) or OFFLINE (N). If the agent is offline, you can remove it from the Navigator and from a table view of managed system attributes: Right-click the table row and select ✎ Clear offline entry from the pop-up menu; or select the offline ▣ Agent Navigator item from the Navigator and click ✖ **Remove**.

Timestamp The date and time at the Tivoli Enterprise Monitoring Server when the agent heartbeat occurred. The format is MM/DD/YY HH:MM:SS.

Type Indicates whether the node is a monitoring agent (Virtual) or Tivoli Enterprise Monitoring Server (blank).

Version The version number of the monitoring agent. The number is formatted as 01.50.00, where 01 is the major version number, 05 is the secondary version number, and 00 the tertiary.

Situation Definition attributes

The Situation Definition attributes supply you with the complete definition of situations. You can apply the predefined query for this group (in the Query editor under Tivoli Enterprise Monitoring Server → Situation Definition) to charts and tables; you cannot use these attributes in situations.

Advice The expert advice for a situation that appears in the event workspace.

Affinities The internal identification number for the agent and attribute group referenced in the situation.

Alert List Whether or not events for this situation belong to an alert list.

Auto SOPT The auto SOPT.

Auto Start Whether the situation is set to autostart, *YES, or to be started manually, *NO.

Command The type of command specified in the Action tab.

Description The description written for the situation.

Destination Node The managed system to which the situation is distributed.

Event Attribute A single digit to describe the type of event.

HUB The name of the hub monitoring server that controls the situation.

Interval Days A two-digit number to express the interval between data samplings.

Interval Seconds A six-digit number that expresses the time between data samplings.

Last CCSID Internal identifiers of the situation expressions.

Last Date The date of the last revision to the situation. The format is CYYMMDDHHMMSSmmm (as in 1050415074501000 for April 15, 2005 at 07:45:01) where:

    C = Century (0 for 20th, 1 for 21st)

    Y = Year

    M = Month

    D = Day

    H = Hour

M = Minute

S = Second

m = millisecond

Last Release The last release to associate with the situation.

Last user The user ID of the last person to update the situation.

Location Flag The location flag identifier.

NOCOL Whether this is a no column attribute.

Notify Arguments Any notification arguments.

Notify Options Any notification options.

Object Lock The object lock.

Predicate The formula, including the detailed names of the attribute group and attributes. .

QIBCLASSID The query information base class ID.

QIBSCOPE The scope of the query information base.

Reflex OK Whether an action command has been set.

Send Message Q The universal message that will be sent when the situation is true.

Sit Info Any other information about the situation.

Situation Name The unique identifier given to the situation or policy.

Full Name The name given to the situation or policy. The name can be up to 256 bytes long.

Source The situation source.

Situation State Event attributes

The Situation Event attributes collect status information on situation events. Views that display this information include My Acknowledged Events, Selected Event, and Similar Events and comprise some of the predefined workspaces for the Enterprise Navigator item.

Affinities The affinity identifier.

Close Indicator The close indicator for conveying how the event was closed, such as closed by a user or by a process and whether the event record contains notes.

Close Time Not in use. The time at the hub monitoring server when the event was closed.

Display Item For multiple-row attribute groups, enables the situation to continue sampling rows with a different display item than the one that caused the situation to become true.

Expiration Interval The length of time assigned to keep an acknowledgement open before it expires automatically. Valid values are four digits.

Expiration Time The time that the acknowledgement is scheduled to expire. Valid values follow the MM/DD/YY HH:MM:SS format, as in 12/30/08 23:59:59.

Full Name The name given to the situation.

Group ID The identifier for the group that the situation belongs to.

HUBCMS The hub monitoring server that the portal server connects to.

Impact The name of the Navigator item where the event originated.

Last Update Time The time of the last event activity, either the last user action or a significant system-generated action.

Local Timestamp The time the event occurred or the condition was observed by the originating managed system. The format is MM/DD/YY HH:MM:SS, as in 10/14/08 13:45:35 for October 14, 2008 at 1:45:35 PM.

LocFlag Not in use. The location flag.

Misc Not in use. Miscellaneous event data.

MOBJ The Navigator item the situation is associated with (where the event indicator originates).

Node The name of the monitoring server where the situation is running.

Opened The time at the managed system when the situation status changed.

Owner The user ID of the person who owns this event.

Pathname Not in use. The path to any attachments and saved notes.

Reference ID An identifying attribute to further qualify the event, such as a situation override.

Severity The severity value of the situation event. Valid values:

Severity level	Enumerated value
Fatal	100
Critical	90
Minor	40
Warning	30
Harmless	27
Informational	25
Unknown	15

Situation Name The name given to the situation or policy.

Source The system host name or other name derived from the monitoring agent that identifies the source of the event.

Status The event status, such as Closed or Acknowledged. Valid values:

Event status	Enumerated value
Open	Y
Closed	N
Stopped	P
Started	S
Deleted	D
Acknowledged	A
Reopened	E
Expired	F
Problem	X

TEC Severity The severity value of the event from the Tivoli Enterprise Console.

Transfer Owner The ID of the user who is to receive ownership of the event.

Type Sampled (the situation samples at intervals), pure (unsolicited notification; the situation has no sampling interval), or hub (situation that runs at the hub monitoring server): 0, 1, 2.

UID The event record identifier, consisting of the current record count. After the counter reaches 9999, it starts at 0001 again.

UUID A unique identifier assigned to the event.

Situation Status Current attributes

The Active Situation Status attributes collect status information on situations and policies that are running. The situation event console and message log views display these attributes. You can apply one of the predefined queries for this group (in the Query editor under Tivoli Enterprise Monitoring Server → Situation Status Current) to charts and tables; you cannot use these attributes in situations.

Display Item Attribute to further qualify the situation. With a display item set for a multiple-row attribute group, the situation continues to look at the other rows in the sampling and opens more events if other rows qualify. This optional addition to a situation also enables you to quickly distinguish events in the message log and event console views. The display item is designated in the Situation editor (Formula tab, click **Advanced**).

Full Name The name given to the situation or policy.

Global Timestamp The time the event was recorded by the hub Tivoli Enterprise Monitoring Server. The format is MM/DD/YY HH:MM:SS.

Local Timestamp The time the event occurred or the condition was observed by the originating managed system. The format is MM/DD/YY HH:MM:SS, as in 07/23/09 04:45:35 for July 23,2009 at 4:45:35 AM.

Name The unique identifier given to the situation or policy.

Node The name of the monitoring server where the situation is running.

Origin Node The name of the managed system where the situation is being monitored.

Owner The user ID of the person who acknowledged the event.

Sound File The name of the sound file associated with the Navigator item for this situation.

Status The status of the situation. The possible values are as follows:

Acknowledged	A	The situation event has been acknowledged.
Deleted	D	The situation has been deleted.
Problem	X	The situation is in a problem state.
Expired	F	The acknowldedgement has expired and the situation is still true.
Open	Y	The situation is running and is true.
Closed	N	The situation is running, has been true, and is now false.
Reopened	E	The acknowledgement was removed before it had expired and the situation is still true.
Started	S	The situation is being started.
Stopped	P	The situation has been stopped.

Type The event is either pure (unsolicited notification; the situation has no sampling interval) or sampled (the situation is sampled at intervals).

Sampled

> Sampled events happen when a situation becomes true. Situations sample data at regular intervals. When the situation is true, it causes an event, which is closed automatically when the situation goes back to false (or you can close it manually if you stop the situation).

Pure

> Pure events are unsolicited notifications. Examples of pure events are an out-of-paper condition on a printer and an occurrence that adds an entry to the Windows System Log. Some monitoring agents have attributes that report pure events, such as the Windows Event Log and Windows File

Change attribute groups. A situation using one of these attributes can monitor for pure events. Because of the nature of pure events and the fact that a situation created with pure event attributes does not have a sampling interval, they are not closed automatically like sampled events; you must close the event manually or through an UNTIL clause.

Note: If you create a query and apply a column function to the Type attribute, pure events will not be included in the AVG, COUNT, MAX, MIN or SUM calculation.

Situation Status Log attributes

The Situation Status attributes collect status information on started situations and policies. You can apply any of the predefined queries for this group (in the Query editor under Tivoli Enterprise Monitoring Server → Situation Status Log) to charts and tables; you cannot use these attributes in situations.

Display Item Attribute to further qualify the situation. With a display item set for a multiple-row attribute group, the situation continues to look at the other rows in the sampling and opens more events if other rows qualify. This optional addition to a situation also enables you to quickly distinguish events in the message log and event console views. The display item is designated in the Situation editor (Formula tab, click **Advanced**).

Full Name The name given to the situation or policy.

Global Timestamp The time the event was recorded by the hub Tivoli Enterprise Monitoring Server. The format is MM/DD/YY HH:MM:SS.

Local Timestamp The time the event occurred or the condition was observed by the originating managed system. The format is MM/DD/YY HH:MM:SS, as in 07/23/09 04:45:35 for July 23,2009 at 4:45:35 AM.

Name The unique identifier given to the situation or policy.

Node The name of the monitoring server where the situation is running.

Origin Node The name of the managed system where the situation is being monitored.

Owner The user ID of the person who acknowledged the event.

Sound File The name of the sound file associated with the Navigator item for this situation.

Status The status of the situation. The possible values are as follows:

Acknowledged	A	The situation event has been acknowledged.
Deleted	D	The situation has been deleted.
Problem	X	The situation is in a problem state.
Expired	F	The acknowledgement has expired and the situation is still true.

⬜ Open	Y	The situation is running and is true.
☑ Closed	N	The situation is running, has been true, and is now false.
🔄 Reopened	E	The acknowledgement was removed before it had expired and the situation is still true.
▶ Started	L	The situation is being started.
🔴 Stopped	P	The situation has been stopped.
Not Available	<null>	The status for the situation is not being returned. To investigate, set this client trace: `error (unit:Managed detail)`

To determine why the status is set as it is, the following client trace will help: `error (unit:Managed detail)`.

Type The event is either pure (unsolicited notification; the situation has no sampling interval) or sampled (the situation is sampled at intervals).

Sampled

> Sampled events happen when a situation becomes true. Situations sample data at regular intervals. When the situation is true, it causes an event, which gets closed automatically when the situation goes back to false (or you can close it manually if you stop the situation).

Pure Pure events are unsolicited notifications. Examples of pure events are an out-of-paper condition on a printer and an occurrence that adds an entry to the Windows System Log. Some monitoring agents have attributes that report pure events, such as the Windows Event Log and Windows File Change attribute groups. A situation using one of these attributes can monitor for pure events. Because of the nature of pure events and the fact that a situation created with pure event attributes does not have a sampling interval, they are not closed automatically like sampled events; you must close the event manually or through an UNTIL clause.

> **Note**: If you create a query and apply a column function to the Type attribute, pure events will not be included in the AVG, COUNT, MAX, MIN or SUM calculation.

Universal Messages attributes

Universal Messages attributes keep track of the messages that have been generated through a situation or policy action command. The Universal Message Console view uses these attributes, and they can be used to create situations.

Category The category name that was given for the message, such as 'Information' or 'SiteCritical'.

Message Severity The severity name that was given for the message to indicate its importance, such as 'Low' or '5'.

Message Text The content of the message, up to 256 characters. You can use the **SCAN** function to find a value anywhere in the message, and it can include an asterisk (*) as a wildcard, such as inciden* to find messages containing incidence, incident, or incidental. You can also use the **STR**ing function to look for matching text at a specific location in the message.

Originnode The network name of the managed system where the situation triggering the Universal Message became true. Valid entries are up to 32 characters.

Situation Name The name of the situation generating the Universal Message.

Timestamp The date and time at the Tivoli Enterprise Monitoring Server when the universal message was sent. The format is MM/DD/YY HH:MM:SS.

Universal Time attributes

The Universal Time attributes show the time and date at the monitoring agent expressed in UTC (Coordinated Universal Time) format. For example, if the UTC time is 19:15.00 for an agent in New York, the local time is 2:15 PM Standard Time. These attributes are useful for action situations triggered by time.

Note:
1. When building situations, you will see that the Universal Time attribute group can be combined with any attribute group to create a condition.
2. The attributes are not available in Tivoli Enterprise Monitoring Server queries.
3. Seconds, Minutes, Hours, and Day Of Month are treated as two-character text attributes. For example, == 05 Minutes is correct, but == 5 Minutes will never be true.

Day Of Month The day in the month at the agent location when the data was sampled, from 01 to 31. The first day of the month is 01.

Day Of Week The numeric representation for the day in the week at the agent location when the data was sampled, from Sunday to Saturday. Select the day from the ▾ list.

Hours The hour in the day at the agent location when the data was sampled, from 00 to 23, such as 01 for 1 am.

Minutes The minute in the hour at the agent location when the data was sampled, from 00 to 59, such as 05 for the fifth minute.

Month Of Year The numeric representation of the month in the year at the agent location when the data was sampled. Select the month from the ▾ list.

Seconds The seconds into the minute at the agent location when the data was sampled, from 00 to 59, such as 05 for the fifth second.

Time The time at the agent location when the data was sampled, formatted as HHMMSS. For example, 153000 is 3:30 PM

Timestamp The time stamp at the agent location when the data was sampled. The format is MM/DD/YY HH:MM:SS.

Documentation library

Various publications are relevant to the use of IBM Tivoli Monitoring and to the commonly shared components of Tivoli Management Services.

These publications are listed in the following categories:
- IBM Tivoli Monitoring library
- Related publications

Documentation is delivered in the IBM Tivoli Monitoring and OMEGAMON XE Information Center at http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp and also in the **Files** section of the Application Performance Management community.

For information about accessing and using the publications, select IBM Tivoli Monitoring ⇒ **Using the publications** in the **Contents** pane of the IBM Tivoli Monitoring and OMEGAMON XE Information Center at http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp.

To find a list of new and changed publications, click the **New in this release** topic on the IBM Tivoli Monitoring welcome page. To find publications from the previous version of a product, click **Previous versions** under the name of the product in the **Contents** pane.

IBM Tivoli Monitoring library

The IBM Tivoli Monitoring library provides information about the commonly shared components of Tivoli Management Services.

- *Quick Start Guide*

  Introduces the components of IBM Tivoli Monitoring.
- *Installation and Setup Guide*, SC22-5445

  Provides instructions for installing and configuring IBM Tivoli Monitoring components on Windows, Linux, and UNIX systems.
- *High Availability Guide for Distributed Systems*, SC22-5455

  Gives instructions for several methods of ensuring the availability of the IBM Tivoli Monitoring components.
- *Program Directory for IBM Tivoli Management Services on z/OS*, GI11-4105

  Gives instructions for the SMP/E installation of the Tivoli Management Services components on z/OS.
- *Administrator's Guide*, SC22-5446

  Describes the support tasks and functions required for the Tivoli Enterprise Portal Server and clients, including Tivoli Enterprise Portal user administration.
- *Command Reference* available on Service Management Connect

  Provides detailed syntax and parameter information, as well as examples, for the commands you can use in IBM Tivoli Monitoring.
- *Messages* available on Service Management Connect

  Lists and explains messages generated by all IBM Tivoli Monitoring components and by z/OS-based Tivoli Management Services components (such as Tivoli Enterprise Monitoring Server on z/OS and TMS:Engine).

- *Troubleshooting Guide* available on Service Management Connect

  Provides information to help you troubleshoot problems with the software.
- *Tivoli Enterprise Portal User's Guide* available on Service Management Connect

  Complements the Tivoli Enterprise Portal online help. The guide provides hands-on lessons and detailed instructions for all Tivoli Enterprise Portal features.
- Tivoli Enterprise Portal online help

  Provides context-sensitive reference information about all features and customization options of the Tivoli Enterprise Portal. Also gives instructions for using and administering the Tivoli Enterprise Portal.

Documentation for the base agents

If you purchased IBM Tivoli Monitoring as a product, you received a set of base monitoring agents as part of the product. If you purchased a monitoring agent product (for example, an OMEGAMON XE product) that includes the commonly shared components of Tivoli Management Services, you did not receive the base agents.

The following publications provide information about using the base agents.
- Agentless operating system monitors
  - *Agentless Monitoring for Windows Operating Systems User's Guide*, SC23-9765
  - *Agentless Monitoring for AIX Operating Systems User's Guide*, SC23-9761
  - *Agentless Monitoring for HP-UX Operating Systems User's Guide*, SC23-9763
  - *Agentless Monitoring for Solaris Operating Systems User's Guide*, SC23-9764
  - *Agentless Monitoring for Linux Operating Systems User's Guide*, SC23-9762
- OS agent documentation is delivered in the following locations:

  Agent Installation and Configuration Guide
  Available in the Information Center:
  - *IBM i OS Agent Installation and Configuration Guide*, SC27-5653
  - *Linux OS Agent Installation and Configuration Guide*, SC27-5652
  - *UNIX OS Agent Installation and Configuration Guide*, SC27-5651
  - *Windows OS Agent Installation and Configuration Guide*, SC27-5650

  Agent Reference
  Available on Service Management Connect

  Agent Troubleshooting Guide
  Available on Service Management Connect

  Infrastructure Management Dashboards for Servers Reference
  Available on Service Management Connect
- Warehouse agent documentation is delivered in the following locations:

  Agent Installation and Configuration Guide
  Available in the Information Center:
  - *Warehouse Proxy Agent Installation and Configuration Guide*, SC27-5655
  - *Warehouse Summarization and Pruning Agent Installation and Configuration Guide*, SC27-5654

  Agent Reference
  Available on Service Management Connect

Agent Troubleshooting Guide
Available on Service Management Connect

- System P agents
  - *AIX Premium Agent User's Guide*, SA23-2237
  - *CEC Base Agent User's Guide*, SC23-5239
  - *HMC Base Agent User's Guide*, SA23-2239
  - *VIOS Premium Agent User's Guide*, SA23-2238
- Other base agents
  - *Agent Builder User's Guide*, SC32-1921
  - *Performance Analyzer User's Guide*, SC27-4004
  - *Systems Director base Agent User's Guide*, SC27-2872
  - *Tivoli Log File Agent User's Guide*, SC14-7484
  - *Tivoli zEnterprise Monitoring Agent User's Guide*, SC14-7359 and the *Tivoli zEnterprise Monitoring Agent Installation and Configuration Guide*, SC14-7358

Related publications

For information about related products and publications select **OMEGAMON XE shared publications** or other entries in the **Contents** pane of the IBM Tivoli Monitoring and OMEGAMON XE Information Center.

You can access the IBM Tivoli Monitoring and OMEGAMON XE Information Center at http://pic.dhe.ibm.com/infocenter/tivihelp/v61r1/index.jsp .

You can also access other information centers at IBM Tivoli Documentation Central (https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Documentation%20Central).

Tivoli Monitoring community on Service Management Connect

Connect, learn, and share with Service Management professionals: product support technical experts who provide their perspectives and expertise.

For information about Tivoli products, see the Application Performance Management community on SMC at IBM Service Management Connect > Application Performance Management (http://www.ibm.com/developerworks/servicemanagement/apm).

For introductory information, see IBM Service Management Connect (http://www.ibm.com/developerworks/servicemanagement).

Use Service Management Connect in the following ways:

- Become involved with transparent development, an ongoing, open engagement between other users and IBM developers of Tivoli products. You can access early designs, sprint demonstrations, product roadmaps, and prerelease code.
- Connect one-on-one with the experts to collaborate and network about Tivoli and the (enter your community name here) community.
- Read blogs to benefit from the expertise and experience of others.
- Use wikis and forums to collaborate with the broader user community.

Other sources of documentation

You can obtain additional technical documentation about monitoring products from other sources.

- Tivoli wikis

  IBM Service Management Connect > Application Performance Management (http://www.ibm.com/developerworks/servicemanagement/apm) includes a list of relevant Tivoli wikis that offer best practices and scenarios for using Tivoli products, white papers contributed by IBM employees, and content created by customers and business partners.

  Two of these wikis are of particular relevance to IBM Tivoli Monitoring:

  – The IBM Tivoli Monitoring Wiki (https://www.ibm.com/developerworks/community/wikis/home?lang=en#!/wiki/Tivoli%20Monitoring) provides information about IBM Tivoli Monitoring and related distributed products, including IBM Tivoli Composite Application Management products.

  – The Tivoli System z® Monitoring and Application Management Wiki provides information about the OMEGAMON XE products, NetView for z/OS, Tivoli Monitoring Agent for z/TPF, and other System z monitoring and application management products.

- IBM Integrated Service Management Library

  http://www.ibm.com/software/brandcatalog/ismlibrary/

  IBM Integrated Service Management Library is an online catalog that contains integration documentation and other downloadable product extensions.

- Redbooks®

  http://www.redbooks.ibm.com/

  IBM Redbooks and Redpapers include information about products from platform and solution perspectives.

- Technotes

  Technotes provide the latest information about known product limitations and workarounds. You can find Technotes through the IBM Software Support Web site at http://www.ibm.com/software/support/.

Support information

If you have a problem with your IBM software, you want to resolve it quickly. IBM provides ways for you to obtain the support you need.

Online

The following sites contain troubleshooting information:

* Go to the IBM Support Portal (http://www.ibm.com/support/entry/portal/software) and follow the instructions.
* Go to IBM Service Management Connect > Application Performance Management (http://www.ibm.com/developerworks/servicemanagement/apm) and select the appropriate wiki.

IBM Support Assistant

The IBM Support Assistant (ISA) is a free local software serviceability workbench that helps you resolve questions and problems with IBM software products. The ISA provides quick access to support-related information and serviceability tools for problem determination. To install the ISA software, go to IBM Support Assistant (http://www-01.ibm.com/software/support/isa).

Troubleshooting Guide

For more information about resolving problems, see the product's Troubleshooting Guide.

Using IBM Support Assistant

The IBM Support Assistant is a free, stand-alone application that you can install on any workstation. You can then enhance the application by installing product-specific plug-in modules for the IBM products you use.

The IBM Support Assistant saves you the time it takes to search the product, support, and educational resources. The IBM Support Assistant helps you gather support information when you need to open a problem management record (PMR), which you can then use to track the problem.

The product-specific plug-in modules provide you with the following resources:

* Support links
* Education links
* Ability to submit problem management reports

For more information, and to download the IBM Support Assistant, see http://www.ibm.com/software/support/isa. After you download and install the IBM Support Assistant, follow these steps to install the plug-in for your Tivoli product:

1. Start the IBM Support Assistant application.
2. Select **Updater** on the Welcome page.
3. Select **New Properties and Tools** or select the **New Plug-ins** tab (depending on the version of IBM Support Assistant installed).
4. Under **Tivoli**, select your product, and then click **Install**. Be sure to read the license and description.

If your product is not included on the list under **Tivoli**, no plug-in is available yet for the product.

5. Read the license and description, and click **I agree**.
6. Restart the IBM Support Assistant.

Obtaining fixes

A product fix might be available to resolve your problem. To determine which fixes are available for your Tivoli software product, follow these steps:

1. Go to the IBM Software Support website at http://www.ibm.com/software/ support.
2. Under **Select a brand and/or product**, select **Tivoli**.

   If you click **Go**, the **Search within all of Tivoli support** section is displayed. If you don't click **Go**, you see the **Select a product** section.
3. Select your product and click **Go**.
4. Under **Download**, click the name of a fix to read its description and, optionally, to download it.

   If there is no **Download** heading for your product, supply a search term, error code, or APAR number in the field provided under **Search Support (this product)**, and click **Search**.

For more information about the types of fixes that are available, see the *IBM Software Support Handbook* at http://www14.software.ibm.com/webapp/set2/sas/ f/handbook/home.html.

Receiving weekly support updates

To receive weekly e-mail notifications about fixes and other software support news, follow these steps:

1. Go to the IBM Software Support website at http://www.ibm.com/software/ support.
2. Click **My support** in the far upper-right corner of the page under **Personalized support**.
3. If you have already registered for **My support**, sign in and skip to the next step. If you have not registered, click **register now**. Complete the registration form using your e-mail address as your IBM ID and click **Submit**.
4. The **Edit profile** tab is displayed.
5. In the first list under **Products**, select **Software**. In the second list, select a product category (for example, **Systems and Asset Management**). In the third list, select a product sub-category (for example, **Application Performance & Availability** or **Systems Performance**). A list of applicable products is displayed.
6. Select the products for which you want to receive updates.
7. Click **Add products**.
8. After selecting all products that are of interest to you, click **Subscribe to email** on the **Edit profile** tab.
9. In the **Documents** list, select **Software**.
10. Select **Please send these documents by weekly email**.
11. Update your e-mail address as needed.
12. Select the types of documents you want to receive.
13. Click **Update**.

If you experience problems with the **My support** feature, you can obtain help in one of the following ways:

Online
> Send an e-mail message to erchelp@ca.ibm.com, describing your problem.

By phone
> Call 1-800-IBM-4You (1-800-426-4968).

Contacting IBM Software Support

IBM Software Support provides assistance with product defects. The easiest way to obtain that assistance is to open a PMR or ETR directly from the IBM Support Assistant.

Before contacting IBM Software Support, your company must have an active IBM software maintenance contract, and you must be authorized to submit problems to IBM. The type of software maintenance contract that you need depends on the type of product you have:

- For IBM distributed software products (including, but not limited to, Tivoli, Lotus®, and Rational® products, as well as DB2 and WebSphere products that run on Windows or UNIX operating systems), enroll in Passport Advantage® in one of the following ways:

    Online
    > Go to the Passport Advantage website at http://www-306.ibm.com/ software/howtobuy/passportadvantage/pao_customers.htm .

    By telephone
    > For the telephone number to call in your country, go to the IBM Software Support website at http://techsupport.services.ibm.com/ guides/contacts.html and click the name of your geographic region.

- For customers with Subscription and Support (S & S) contracts, go to the Software Service Request website at https://techsupport.services.ibm.com/ssr/ login.

- For customers with Linux, iSeries®, pSeries, zSeries, and other support agreements, go to the IBM Support Line website at http://www.ibm.com/ services/us/index.wss/so/its/a1000030/dt006.

- For IBM eServer™ software products (including, but not limited to, DB2 and WebSphere products that run in zSeries, pSeries, and iSeries environments), you can purchase a software maintenance agreement by working directly with an IBM sales representative or an IBM Business Partner. For more information about support for eServer software products, go to the IBM Technical Support Advantage website at http://www.ibm.com/servers/eserver/techsupport.html.

If you are not sure what type of software maintenance contract you need, call 1-800-IBMSERV (1-800-426-7378) in the United States. From other countries, go to the contacts page of the *IBM Software Support Handbook* on the web at http://www14.software.ibm.com/webapp/set2/sas/f/handbook/home.html and click the name of your geographic region for telephone numbers of people who provide support for your location.

Notices

This information was developed for products and services offered in the U.S.A. IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not give you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785 U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan, Ltd.
19-21, Nihonbashi-Hakozakicho, Chuo-ku
Tokyo 103-8510, Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law :

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement might not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Corporation
2Z4A/101
11400 Burnet Road
Austin, TX   78758   U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurement may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to

IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. You may copy, modify, and distribute these sample programs in any form without payment to IBM for the purposes of developing, using, marketing, or distributing application programs conforming to IBM's application programming interfaces.

Each copy or any portion of these sample programs or any derivative work, must include a copyright notice as follows: © (your company name) (year). Portions of this code are derived from IBM Corp. Sample Programs. © Copyright IBM Corp. 2013. All rights reserved.

If you are viewing this information in softcopy form, the photographs and color illustrations might not be displayed.

Trademarks

IBM, the IBM logo, and ibm.com® are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe and all Adobe-based trademarks are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, other countries, or both.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, and Windows are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Privacy policy considerations

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable

information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

Depending upon the configurations deployed, this Software Offering may use session cookies that collect each user's user name for purposes of session management, authentication, and single sign-on configuration. These cookies cannot be disabled.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at http://www.ibm.com/privacy and IBM's Online Privacy Statement at http://www.ibm.com/privacy/details the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at http://www.ibm.com/software/info/product-privacy.

Index

A

about this document vii
acknowledgement
    open an existing 60
administer users 374
administration
    mode 69
advanced filter 97
agent not connected to the monitoring
    server 322
Agent Operations Log
    query 212
Agent Operations Log workspace 393
agents 9
    deploy depot workspace 401
    deployment status workspace 401,
        402, 404, 405
    deployment workspaces 400
all managed systems situations 267
AND expression 418, 419
application support events 30
architecture 3
area chart 112
    plot area 117
    thresholds 119
    value and category axes 125
ASCII and non-ASCII
    graphic view icon file names 150
    paste into formula editor 419
    threshold icon file names 102
    user ID 370
associate a situation 287
associate situations 255
attribute groups
    pure or sampled 417
    single-row or multiple-row 417
attributes 431
audio files 273
Audit Log workspace 394
automation
    policies 327
Average of group 428
AVG formula function 428

B

background information 327
bar chart 109
    overlay 122
    plot area 117
    thresholds 119
    value and category axes 125
baselines
    add historical 132
    add monitored 126
    add statistics to a chart view 130
    monitored chart examples 128
BETWEEN formula function 420
Boolean AND and OR 418
branching OR expressions 418

browser client 8
browser view
    create 166
    find 48
    launching an external web
        browser 168
    link symbols in the URL address 168
    relative URL reference 167
    toolbar 167
built-in functions 185

C

category axis of chart 125
cell functions 420
CHANGE formula function 421
Change in value 421
chart
    area 112
    bar 109
    circular gauge 114
    linear gauge 115
    pie 107
    plot 110
chart views
    comparison 103
charts
    compare situations with 126
    historical baselining in 132
    legend label 123
    model situations with 290
    monitored baselines examples 128
    statistical baselining in 130
    zoom in on an area 46
Check for missing items 422
circular gauge 114
clear a formula editor cell, column, or
    row 419
color 125
column functions 427
common event console view
    create 158
    details 158
Compare date/time 421
Compare to a time +/- delta 425, 426
configure historical data collection 221,
    225
Configure Tivoli Application Dependency
    Discovery Manager Integration 347
console
    common event
        filter, column value 161
        quick filter, removing 162
    event
        quick thresholds, remove 162
        setting view properties 162
cookies 469
copy the contents of a cell or row in the
    formula editor 419
copyright 467
COUNT formula function 428

Count of group members 428
CPU
    high on monitoring server 325
create a new workspace 72
create situation 269
creating a new view 85, 104
custom audio files 273
custom color 125
custom navigator views 237
    add a child item 240
    characteristics 237
    creating a view 238
    delete a view or item 244
    edit a view 242
    edit item properties 243
    renaming a view or item 244
customer support 465
customizing workspace views
    creating a new view 85, 104
    editing a view 86
cut the contents from a formula editor
    cell or row 419

D

Dashboard Application Services Hub
    user IDs 381
DATE formula function 421
delete or restore a workspace 74
delete the contents from a row in the
    formula editor 419
deploy
    depot package list workspace 401
    status summary by transaction
        workspace 402
    status summary workspace 401
developerWorks 462
dissociate a situation 288
distinguished name
    mapping to TEP user ID 368
distribute a historical collection to
    managed systems 228
dynamic thresholding 299

E

editing a view 86
EIB Change Log 213
EIB change log workspace 396
EIB Change Log workspace
    accessing the workspace 396
    what the workspace shows 397
EIF
    map enumerations window 284
    slot customization window 281
EIF receiver
    forward event 280
embedded situation
    edit 419
embedding situations 312

IBM ®

Printed in USA