OPTIMIZE

MERCURY APPLICATION MAPPING™

User's Guide

MERCURY™

BUSINESS TECHNOLOGY OPTIMIZATION

Mercury Application Mapping
User's Guide
Version 6.1

Document Release Date: May 4, 2006

MERCURY™

Mercury Application Mapping User's Guide, Version 6.1

If you have any comments or suggestions regarding this document, please send them by e-mail to documentation@mercury.com.

MAMUG6.1/02

Table of Contents

PART V: GENERATING CONFIGURATION MANAGEMENT AND CHANGE REPORTS

PART XII: APPENDIXES

Welcome to Mercury Application Mapping User's Guide

This guide provides detailed instructions on how to use Mercury Application Mapping.

How This Guide is Organized

The guide contains the following parts:

Part I   **Introduction to Mercury Application Mapping**

Introduces Mercury Application Mapping technology, its architecture, its workflow, and provides a quick tour of the application.

Part II   **Topology Query Language Builder**

Describes how to define and save different types of TQL queries as well as TQL validation restrictions.

Part III   **Pattern Views and Folding Rules**

Describes how to define managed views, which are a series of rules and definitions for displaying query results, and how to notify specific users, group members, or unregistered interested parties of changes that occur in views and nodes.

Part IV   Viewing Queries with Topology View

Describes how to use the to view multi-level maps displaying the results of the query based on a selected view, how to manage the events occurring in the system, at the CI level and at higher, top-view levels, enabling you to locate problems occurring in different areas of your IT infrastructure, and how to monitor the changes occurring in the managed world and provide information about those changes using the Event System.

Part V   Generating Configuration Management and Change Reports

Describes how to generate reports.

Part VI   Correlation Rules

Describes how to use the Correlation Manager to create topology correlations, and to define correlation rules.

Part VII   Enrichment Nodes and Rules

Describes how to use the Enrichment Manager to create Enrichment nodes and rules, and to create sample views.

Part VIII System Reports

Describes how to use the Report Manager to create system reports and how to create a sample system report

Part IX   The CI Type Manager

Describes how to use the CI Type Manager to view the CI Type Model of your managed world, to create and add new CITs to the CI Type Model, and to create a predefined list whose values define an attribute type.

Part X   Mercury Application Mapping Applications

Describes how to use one of the Mercury Application Mapping applications that enables you to take a snapshot of a specific view, save it and then compare it, using Mercury Application Mapping's comparison capabilities, to snapshots taken of that view at different times, how to use another application that enables you to view the content of configuration files and compare the differences between them, and how to assess the impact of infrastructure events.

Part XI   Mercury Application Mapping WebView

Describes how to generate system reports and view them in Mercury Application Mapping WebView (a web-based management console designed to enable business service management).

Part XII   Appendixes

Provides information about the permitted relationships between CIs, key attributes of CITs, relationship definitions and a glossary.

Getting More Information

For information on using and updating the Mercury Application Mapping Documentation Set, reference information on additional documentation resources, typographical conventions used in the Documentation Set, and quick reference information on deploying, administering, and using Mercury Application Mapping, refer to *Getting Started with Mercury Application Mapping*.

Part I

Introduction

1

Introduction to Mercury Application Mapping

This chapter provides a general introduction to Mercury Application Mapping.

About Mercury Application Mapping

Mercury Application Mapping is an application that uses Mercury's patented pending Topology Information System (TIS) technology to manage all the Configuration items (CIs) contained in a managed world. A managed world refers to any self-contained environment that can be described using a topology model. For example, the IT infrastructure of a large business represents a managed world, where the topology is comprised of multiple layers such as networks, protocols, databases, operating systems, and so forth.

Mercury Application Mapping addresses the following operational and functional needs:

➤ **IT resources and application alignment** – automatic discovery of IT resources and their interdependencies from a business service perspective

➤ **Problem isolation and problem resolution** – understanding the casual relations between CIs to locate and address the root-cause of infrastructure problems and reduce troubleshooting time

➤ **Business impact analysis** – evaluating the impact of infrastructure events on business services to assess their implications and prioritize the responses to them

➤ **Asset and change management control** – automatic detection of infrastructure changes, which leads to automatic updating of all the relevant sub-systems

➤ **Performance management and capacity planning**

➤ **Architecture and infrastructure planning**

➤ **Customized category management, such as performance and change**

Mercury Application Mapping offers a configurable and flexible approach to business service management. It uses TQL query language, which enables you to focus on the information relevant to your needs (for details, see "Mercury Application Mapping Topology Query System" on page 9). Mercury Application Mapping enables you to manage views so that you view the information in exactly the format you require For details about managed views, see Chapter 7, "Working with Pattern Views.". Furthermore, the information contained in the results of each query is updated automatically with the latest data entering the Configuration Management Database (CMDB). As a result, once a query and managed view have been defined, they continue to provide updated information about the current state of your managed world. Managed views are displayed in multi-level maps that enable you to identify key CIs and events, as required. The Report Manager enables you to create reports (in HTML, Excel or Table format) about information collected by the system.

Mercury Application Mapping System Architecture

The following diagram provides a graphical overview of the Mercury Application Mapping system architecture:
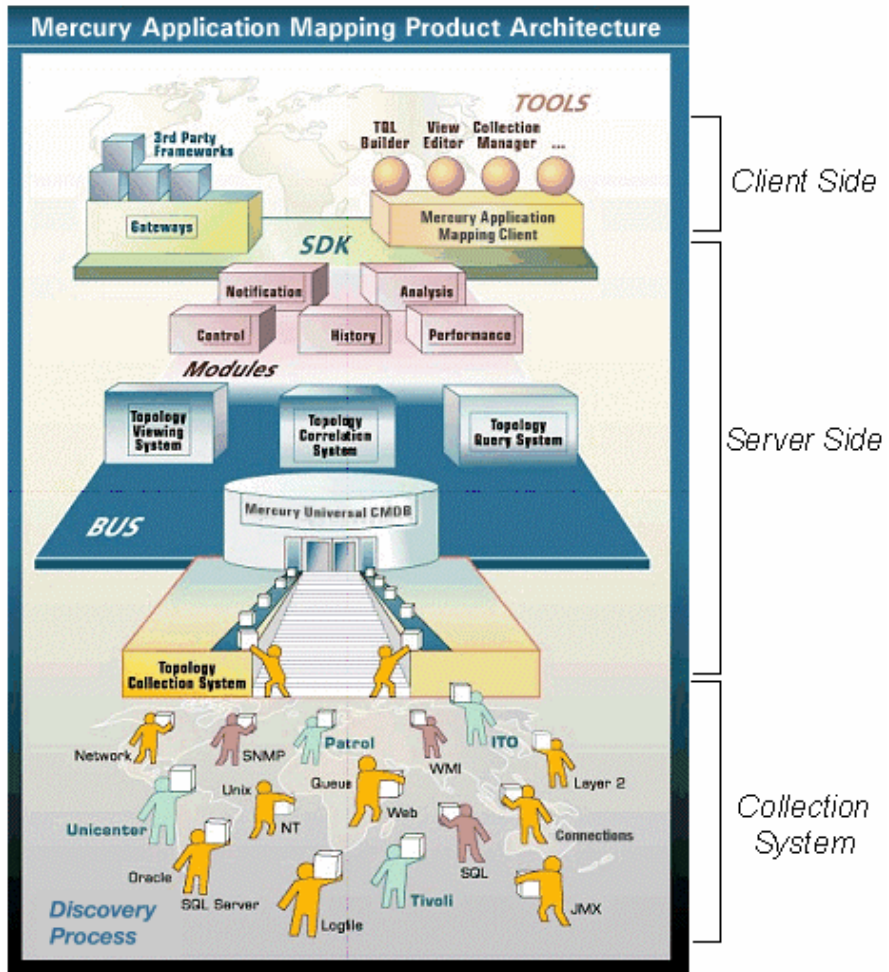
Mercury Application Mapping Discovery Process

Mercury Application Mapping's discovery system enables you to automatically discover and map logical application assets in Layers 2 through 7 of the Open System Interconnection (OSI) Model. It is an ongoing, automatic process that continuously detects changes that occur in your IT infrastructure and updates the CMDB accordingly. It does not require any agent installation on the discovered device and uses a spiral discovery model that has over 100 out-of-the-box discovery patterns.

The discovery process is a gradual uncovering of the elements in your system. Discovery is first done at the most basic level, and then at more in-depth ones. Every time a discovery pattern is activated, it discovers more CIs, which in turn are used as triggers for other discovery patterns. This process continues until your entire IT infrastructure is discovered and mapped.

The discovery process does the following:

➤ Sweeps the application domain

➤ Identifies application element details

➤ Determines which ports are being used

➤ Checks the OS (operating system)

➤ Communicates with Industry standard or Application APIs

For more information on the discovery process, see the *Mercury Application Mapping Discovery Process Tutorial*.

Discovery Process Architecture

The architecture of the Mercury Application Mapping discovery process consists of the following components:

➤ "Discovery Pattern" on page 8

➤ "Discovery Module" on page 8

➤ "Discovery Probe" on page 8

➤ "Mercury Application Mapping Server" on page 8

Discovery Pattern

The discovery process is run by activating discovery patterns. A discovery pattern is an XML template that contains definitions of one or more discovery tasks. It defines which data is discovered, through which protocol, when to pass the data on to the CMDB, and the executing schedule. A discovery pattern contains a description of the configuration items (CIs) and relationships that are created with each specific pattern. The CI and relationship definitions are taken from the CI Type Model, which contains the definitions of all the configuration item types (CITs) and relationships defined in the system.

Discovery Module

A discovery module contains one or more discovery patterns that together, allow a discovery of a specific technology/application. For example: discovery of the network includes patterns that discover: IPs, servers, switches, routers, and so forth.

Discovery Probe

The Discovery Probe is the main discovery component responsible for receiving discovery tasks from the Mercury Application Mapping server, dispatching them to the IT components and sending the results back to the CMDB through the server.

Mercury Application Mapping Server

The Mercury Application Mapping server delivers discovery requests to the Discovery Probe, receives the discovery results, and stores the collected data in the CMDB.

Discovery Domain

The discovery domain maps a range of network addresses to a specific installed Discovery Probe. The connection between the domain name to the physical installed Discovery Probe is made in the **appilog-remote.properties** file, which is located in: **\<Mercury Application Mapping root directory>\root\lib\collectors**.

Discovery Protocol

Discovery of the different IT infrastructure components is performed by using various protocols such as SNMP, WMI, JMX, Telnet, and so forth.

Mercury Application Mapping Topology Query System

The Topology Query System includes Topology Query Language (TQL) that is a language and a tool for discovering, organizing and managing IT infrastructure data (for details, see "Topology Query Language (TQL)" on page 10).

TQL is used to build a business-service model that guides the discovery and identification of the business services and structures the way they are organized and managed in the CMDB. TQL is used to create queries that retrieve business service data from the CMDB, and display that data. TQL queries constantly search the CMDB for changes that occur in the state of managed resources and inform and update the relevant subsystems (for details, see "The Roles of TQL" on page 10).

The Mercury Universal CMDB is the core information repository of Mercury Application Mapping. It contains the CIT Model and the custom-tailored business-service model, and stores and handles the infrastructure data collected and updated by the discovery process (for details, see "Mercury Universal CMDB" on page 11).

The Topology Query System is the underlying mechanism of all Mercury Application Mapping core modules. It enables the user to create a TQL query (for details, see "Mercury Application Mapping Core Modules" on page 12).

Topology Query Language (TQL)

TQL extends the standard SQL language by adding two important capabilities:

➤ TQL enables Mercury Application Mapping to draw conceptual relationships between CIs, which represent their actual interdependencies. Using pre-defined operators, the different types of interconnections that exist between CIs can be established, and consequently the infrastructure design and performance is more accurately represented. This representation serves as a basis and a model for the discovery, arrangement, query and management of complex infrastructures.

➤ TQL has a graphical aspect, which consists of visual symbols and syntax that represent the resources and their interconnections. This visualization of an IT infrastructure, simplifies the understanding, monitoring and managing of the IT business operations.

The Roles of TQL

TQL is not in itself a stand-alone product. It is an integral part of Mercury Application Mapping and it plays several roles:

➤ Builds a business-service model that defines and delineates the interconnection between IT assets, which function in concert as business services. The business-service model guides the discovery and identification of these business services, from the ever-increasing number and complexity of infrastructure resources. Once the resources that comprise the business services are discovered, the business-service model structures the way they are organized and managed in the CMDB.

➤ Creates queries that retrieve business service data from the CMDB, and displays it in a visual representation that facilitates the data monitoring and managing.

➤ Constantly searches the CMDB for changes that occur in the state of managed resources. When such changes are detected, the relevant subsystems are informed and updated.

Mercury Universal CMDB

The Mercury Universal CMDB is the core information repository of Mercury Application Mapping. It contains the CIT Model and the custom-tailored business-service model, and stores and handles the infrastructure data collected and updated by the discovery process. The information pertaining to discovered CIs and relationships is deposited, grouped and updated in the CMDB in the form of CIT definitions according to CI-oriented methodology.

All the data that is stored in the CMDB can be represented in a visual image, which is composed of nodes and relationships, as shown in the following figure:
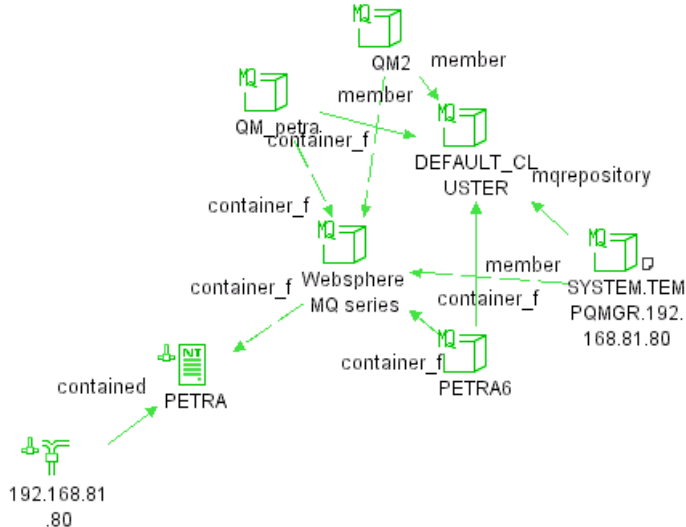
The nodes and relationships are organized in the CMDB according to CI-oriented methodology.

Mercury Application Mapping Core Modules

The Topology Query System is the underlying mechanism of all Mercury Application Mapping core modules. It enables you to create a TQL query, a graphic scenario of CIs and their interdependencies, which addresses the corresponding structure of a business service or any other formation of IT data in the managed world.

After a TQL query is defined, it scans the CMDB for the requested information on a regular basis. When matching results are found, they are stored in one of the database subsets. Then, they can be displayed and analyzed through the other Topology modules.

Since the TQL queries are self-maintaining and dynamic, they are constantly polling the CMDB for changes in CIs or their interdependencies. When changes are discovered, the query results are automatically updated and all relevant subsystems are informed and alerted.

There are five types of TQL queries: Discovery, Viewing, Correlation, Reporting and Enrichment nodes:
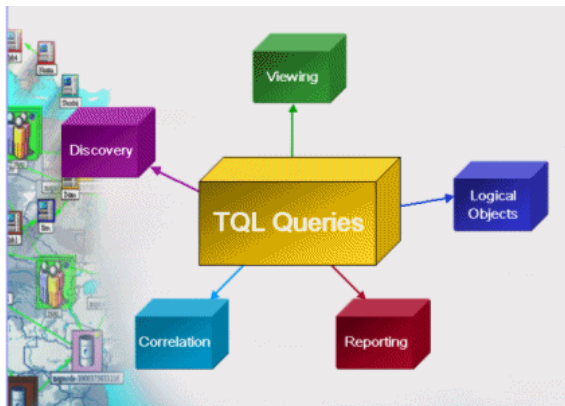
➤ **Viewing System (Mapping System)** – Generates specific managed views (maps) that display the query results and enable their monitoring and management. The Viewing System provides the user with three sets of rules for achieving the kind of display that serves his/her needs best:

➤ **Organizational Rules** – For arranging the results data in a structured display. The organization rules arrange the result data in multi-layer views, and define the position of CIs in each layer. This way, all the query results can fit into one window, and the user can focus on a specific area of interest by navigating between view's layers, either by individually selecting the required layers or by drilling down from higher layers to lower ones. For details, see "Adding Organization (Folding) Rules to Relationships" on page 150.

➤ **Personalization Rules** – For customizing the built-in display rules of the Viewing System. These built-in rules are designed to create visual indications that assist in the displayed data management. By customizing them, the user can adapt the display rules to management requirements. For example, a Viewing System rule states that when an event with severity higher than Normal occurs in relation to a CI in a Normal state, the CI starts blinking and this blinking passes on to the layers above the CI. The user can control this blink propagation and decide whether the blinking propagates and to which view layers. For details, see "Understanding the Event System" on page 234.

➤ **Correlation System** – Defines casual relationship between CIs, which enables to locate and identify, in case of system's alerts, the CI that is the root-cause of a problem and the one(s) that is affected by it.

The Correlation Module uses topology knowledge that is gathered through TQL queries, to automatically outline the interdependencies that exist between service components in a specific IT infrastructure. By using this information to predefine correlation rules and scenarios, correlation events are generated once a change in one CI's state affects the state of others. These automatically-generated correlation events contain information that enable Mercury Application Mapping to quickly identify the true origin of a change, to assess the implication and consequences of infrastructure changes, and to appropriate prioritization in handling alerts. For details, see "Using the Correlation Manager" on page 277.

➤ **Reporting System** – Defines templates for displaying and analyzing the data that is gathered by TQL queries. These templates are defined by the user, which can determine the content, structure and layout of the final reports. The wide variety of options enables the generation of focused summaries and statistics about specific aspects or general conditions of the managed world. For details, see "Creating System Report Templates" on page 347.

➤ **Enrichment System** – Enlarging the CMDB, by adding CIs and relationships that are currently not included in it. This system allows you to integrate prior knowledge or knowledge you deduced from accumulated information, with data that is found through the discovery process. For details, see "Defining Enrichment Nodes and Rules" on page 305.

Mercury Application Mapping Deployment

The following diagram provides a graphical overview of the way the Mercury Application Mapping system is typically deployed:
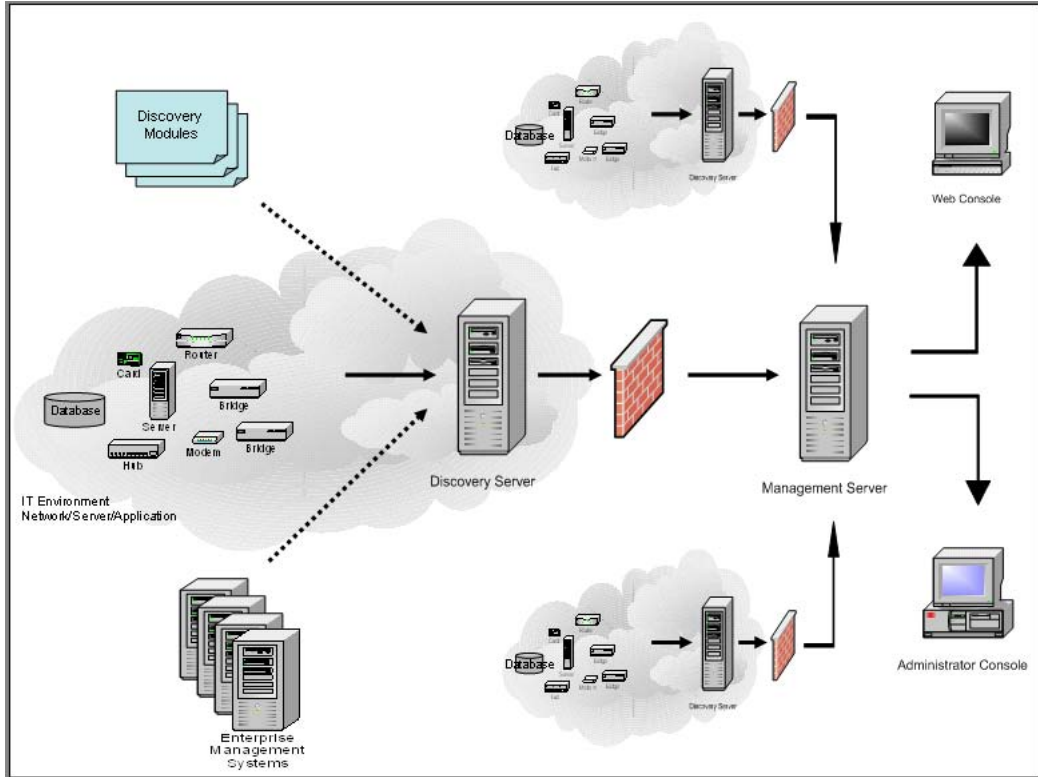
2

Mercury Application Mapping Quick Tour

This chapter introduces the basic concepts of the Mercury Application Mapping application and familiarizes you with its components and functionality.

Mercury Application Mapping Workflow

The following workflow contains the main tasks that are performed when creating a query in Mercury Application Mapping.
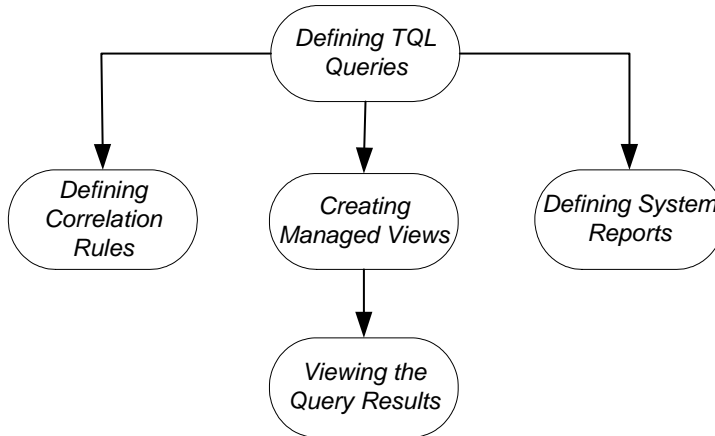
Each step in the workflow summarizes the main tasks performed, as described below:

➤ **Define a new query** – You define a new query, using the Topology Query Language Builder, which consists of selecting TQL nodes and defining the specific attributes and conditions of the nodes and relationships that define the connection between the nodes. Depending on the type of TQL query created, you then create managed views, define correlations rules, or produce reports. For details, see Part II, "Topology Query Language Builder."

➤ **Create a new view** – You create a new view, attach it to a specific View TQL query and define how the query results are displayed in the topology map. This includes adding organizational rules and notifications to define the organization structure of the CIs in the topology map. For details, see Part III, "Pattern Views and Folding Rules."

➤ **View multi-level maps that display the results of queries based on selected views –** You can also search for CIs and display and monitor the changes that occur in your managed data using the Event system. This includes managing the events related to specific nodes in your query and locating problems occurring in your network. For details, see Part IV, "Viewing Queries with Topology View."

➤ **Define the casual relationships between CIs and the implications of these interdependencies** – This consists of creating a Correlation TQL, defining root-cause CI and CIs that are affected by it, deciding which users and/or groups are to receive events from the root-cause CI, the event message format, and the event severity. For details, see Part VI, "Correlation Rules."

➤ **Create and define enrichment rules** – These rules can be used for several purposes: creating new enrichment nodes and relationships for adding data to the CMDB and allowing additional representation options of existing data; updating the value of CI attributes in the database, and deleting CIs from the database. For details, see `

➤ **Create and define system reports based on the Report TQL** – This consists in defining the basic report parameters, the report layout and the data display order. You can then view the selected report or save it to a local directory. For details, see Part VIII, "System Reports."

➤ **Create and Edit CITs –** For details, see Part IX, "The CI Type Manager."

➤ **Manage business services, view event information and perform Mercury Application Mapping actions –** Use Mercury Application Mapping WebView—a Web based management console—to use business service management. For details, see Part XI, "Mercury Application Mapping WebView."

Launching Mercury Application Mapping

The Mercury Application Mapping application is launched using your assigned user name and password. In addition, you must also enter the name of the server that you want to access.

Mercury Application Mapping supplies two default users that are built into the system, as follows:

User Name	Password
Administrator	Administrator
guest	guest

Note: User names and passwords are case-sensitive.

Mercury strongly recommends creating a new administrator–level user name and password after first launching the Mercury Application Mapping application.

To launch Mercury Application Mapping:

1 In the Windows taskbar, select **Start** > **Programs** > **MAM** > **MAM** to open the Login dialog box.

2 Enter your assigned user name and password.

3 In the **Address** box, select the server to which you want to connect.

4 To connect to and manage several servers simultaneously, you can select several servers for display in Mercury Application Mapping by holding down the **Ctrl** key and selecting the required servers, as seen below.
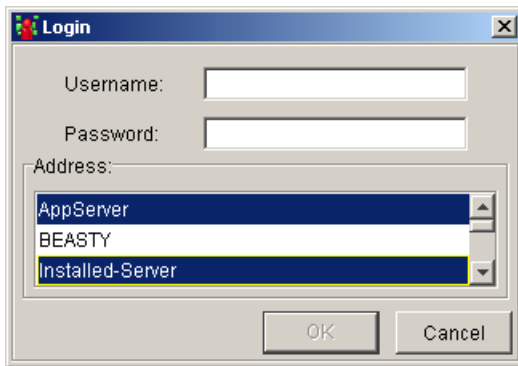
By displaying several servers in one window, you can view and manage them simultaneously and reuse TQLs you prepared for one server for other servers as well.

5 After you enter all the Login details, click **OK** to display the Mercury Application Mapping main window, as described in the next section.

Mercury Application Mapping Main Window

The Mercury Application Mapping main window enables you to work with the Mercury Application Mapping layers, menus, toolbar functions, and windows.

Tool Bar
Menu Bar
Title Bar

View Explorer pane   Topology Map pane   Information pane

Filtered By pane
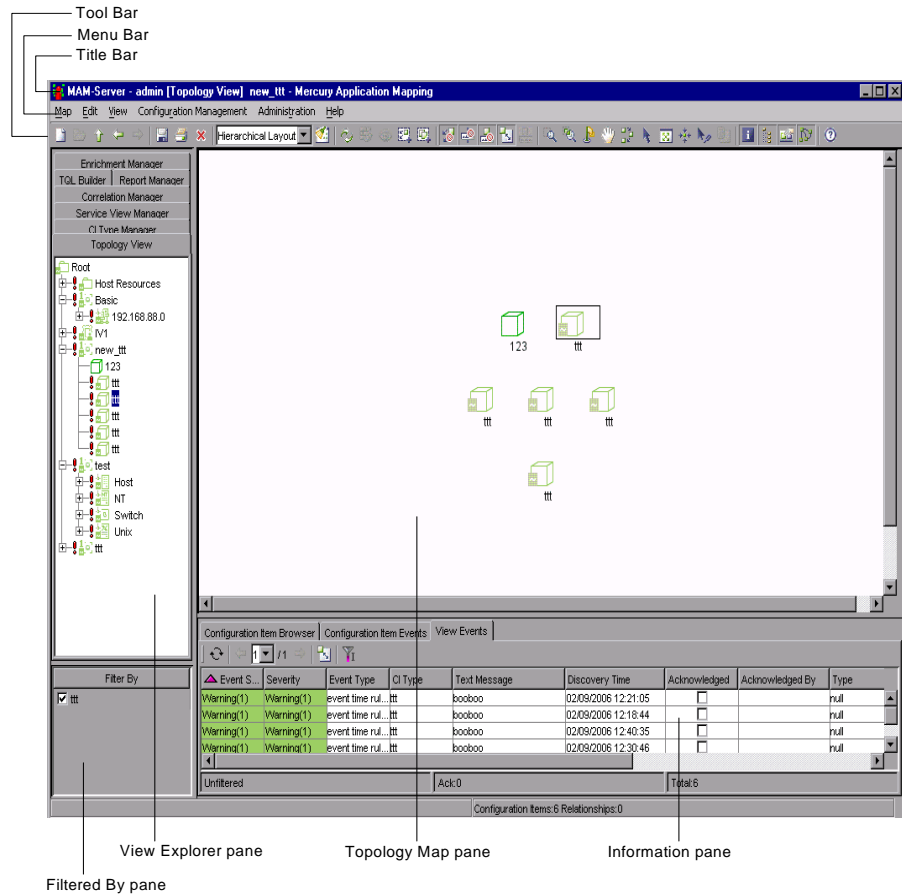
➤ **Title bar** – Displays the following information:

➤ The name of the active server

➤ The name of the currently logged-in user

> ➤ The displayed Manager

> ➤ The name of the selected view/TQL/rule.

➤ **View Explorer pane** – Contains hierarchical tree structures of the CIT Model, TQL Queries, managed views or map CIs, depending on the tab selected (for details, see "View Explorer" on page 23).

➤ **Filtered By pane** – Contains CI groups that can be displayed or hidden in the topology map.

➤ **Topology map** – Displays multi-layered views of CIs and relationships, including the connections between them, using various relationship types. Each CIT has a unique icon; its neighbors (as indicated by its connection to the CIT in the View Explorer) all share the same shape surrounding the icon. (Shapes are predefined according to logical CIT groupings.)

In addition, the topology map includes the following components:

> ➤ When the TQL Builder tab is selected, the Configuration Item Types pane is displayed.

> ➤ When the Topology View tab is selected, the Information pane is displayed.

➤ **Information pane** – Displays different tabs (Information Browsers) depending on the Information Browser selected **View** menu. To hide/display the Information Browsers, select **View** > **Information Browsers** and click the Information Browser you want to hide/display.

➤ **Menu bar –** For details, see "Menu Bar" on page 24.

➤ **Toolbar –** For details, see "Toolbar" on page 34.

This section contains the following topics:

➤ "View Explorer" on page 23

➤ "Menu Bar" on page 24

➤ "Toolbar" on page 34

View Explorer

The View Explorer contains hierarchical tree structures of the CIT Model, TQL Queries, managed views or map CIs, depending on the tab selected. The tabs are:

➤ **Correlation Manager** – Enables you to define correlation rules, which specify a casual relationship between two or more nodes and physical relationships. Each correlation rule you define is attached to a specific Correlation TQL query. This type of TQL query is created especially for correlation rules, and it provides them with the required dependency information for their execution. For details, see "Using the Correlation Manager" on page 277.

➤ **Service View Manager** - Enables you to define managed views that are a series of rules and definitions for displaying query results. For details, see "Introduction to the Service View Manager" on page 99.

➤ **Topology View** – Enables you to view the results of your TQL queries based on the view defined with the Service View Manager. The maps displayed in the topology map may be displayed in multiple layers, depending on whether organization rules are defined in the view. You can drill down from one layer to the next to see more CIs. The topology map also enables you to view events associated with the CIs displayed in each map. For details, see "Introduction to the Topology View" on page 163.

➤ **CI Type Manager** – Enables you to view the information in the CIT Model, which consists of all the defined CITs in the system and the relationships that define the connection between them. The CIT Model represents the structure of the managed world on which Mercury Application Mapping is used. For details, see "Introduction to the CI Type Manager" on page 377.

➤ **Enrichment Manager** – Enables you to integrate prior knowledge of the infrastructure with data that is found and gathered through the Mercury Application Mapping Discovery system and provides you with greater flexibility in managing queries and designing views. For details, see "Introduction to Enrichment Nodes and Relationships" on page 295.

➤ **TQL Builder -** Enables you to define and save different types of TQL queries. These queries, which use a structured language similar to SQL, perform various functions in the system, such as extracting the information from the Mercury Universal CMDB that is relevant to the user, or creating a report containing specific information. For details, see "Introducing the Topology Query Language (TQL) Builder" on page 49.

➤ **Report Manager** – Enables you to define system reports about selected nodes in the managed world. Each report is based on a specific Report TQL query, which is created especially for it. (If required, you can define multiple reports for each query.) The Report Manager enables you to define exactly what information appears in each report, including the order in which nodes are listed and their placement relative to other nodes, how the information within each node is sorted and the maximum number of rows of information to include for each node. System reports can be displayed in three formats (HTML, Table and Excel), and can either be viewed immediately or saved for later use. For details, see "Introduction to the Report Manager" on page 341.

Menu Bar

This section describes the menu options available from the Mercury Application Mapping menu bar.

Map Menu

The **Map** menu contains the following options:

➤ **New** – Enables you to define new items in all Mercury Application Mapping Managers (apart from the topology map):

➤ New queries in the TQL Builder. For details, see "Creating a TQL Query" on page 59.

➤ New views in the Service View Manager. For details, see "Creating a Pattern View" on page 110.

➤ New correlations rules in the Correlation Manager. For details, see "Creating a Correlation Rule" on page 284.

➤ New enrichment nodes and relationships in the Enrichment Manager. For details, see "Creating Enrichment Nodes and Relationships" on page 310.

➤ New System Reports in the Report Manager. For details, see "Creating a Report" on page 349.

➤ New CITs in the CI Type Manager. For details, see "Creating CITs" on page 384.

➤ **Open** – Enables you to open a selected parent CI and drill down to its lower map layers in the topology map. For details, see Chapter 9, "Using the Topology Map."

➤ **Save** – Enables you to save new and changed items in all Mercury Application Mapping Managers (apart from the CI Type Manager):

➤ TQL queries in the TQL Builder. For details, see "Saving a Query" on page 90.

➤ Managed views in the Service View Manager. For details, see "Saving the Pattern View" on page 159.

➤ Correlations rules in the Correlation Manager. For details, see "Saving the Correlation Rule" on page 291.

➤ Any changes to the manual layout in the topology map. For details, see "Improving Views" on page 197.

➤ System Reports in the Report Manager. For details, see "Viewing and Saving the Generated Report" on page 357.

➤ **Import** – Enables you to import into your system TQLs of different types, as follows:

➤ TQL queries in the TQL Builder. For details, see "Saving a Query" on page 90.

➤ Managed views in the Service View Manager. For details, see "Saving the Pattern View" on page 159.

➤ Correlations rules in the Correlation Manager. For details, see "Saving the Correlation Rule" on page 291.

➤ Report templates (for details, see "Introduction to the Report Manager" on page 341).

➤ Enrichment rules (for details, see "Introduction to Enrichment Nodes and Relationships" on page 295).

You can import numerous items at once, by using the CTRL and SHIFT keys.

➤ **Export All** – Enables you to export a whole set of TQL items at once. The TQL sets you can import are as follows:

  ➤ TQL queries in the TQL Builder. For details, see "Saving a Query" on page 90.

  ➤ Managed views in the Service View Manager. For details, see "Saving the Pattern View" on page 159.

  ➤ Correlations rules in the Correlation Manager. For details, see "Saving the Correlation Rule" on page 291.

  ➤ Report templates (for details, see "Introduction to the Report Manager" on page 341).

  ➤ Enrichment rules (for details, see "Introduction to Enrichment Nodes and Relationships" on page 295).

  For details on exporting a set of TQLs, see "Exporting a Set of TQL Items" on page 91.

➤ **Layer Setup** – Enables you to accurately define the view's layer layout, by assigning positions for the nodes and relationships of the view, to achieve the clearest presentation possible. For details, see "Improving Views" on page 197.

➤ **Export Graph** – Enables you to save a topology map to file. For details, see "Saving a Topology Map" on page 218.

➤ **Change Password.** – Enables you to change your login password.

➤ **Page Setup** – Enables you to define the page setup for printing the contents of the topology map.

➤ **Print** – Enables you to print the contents of the Information pane in the topology map. For details, see "Printing the Topology Map" on page 220.

➤ **Map Printing Settings** – Enables you to set printing options. For details, see "Printing the Topology Map" on page 220.

➤ **Print Map** – Enables you to print the contents of the topology map in the topology map. For details, see "Printing the Topology Map" on page 220.

➤ **Print Preview** – Enables you to preview the topology map before printing.

➤ **Exit** – Closes the application.

Edit Menu

The **Edit** menu contains the following options:

➤ **Add Background Image –** Enables you to add a background map for display and navigational purposes in the topology map. This is used with the Manual Layout option. For details, see "Adding a Background Image to a Topology Map" on page 199.

➤ **Find –** Enables you to search for CIs in the topology map, either in specific views and layers, or in the entire CMDB. The Find option also provides you with different search criteria, through which you can search for CIs according to their CIT, label and/or attributes. For details, see "Searching for CIs" on page 223

➤ **Insert Configuration Item**– Displays the Insert Configuration Item dialog box, enabling the administrator to manually add CIs to the database. For details, see "Adding CIs to the Topology Database" in the *Mercury Application Mapping Administration Guide*.

View Menu

The **View** menu contains the following options:

➤ **Layout** – Contains the following options for displaying the contents of the topology map:
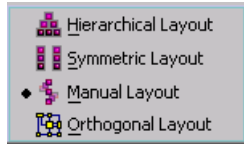
➤ **Hierarchical Layout** – Enables you to display the precedence relations in the topology:
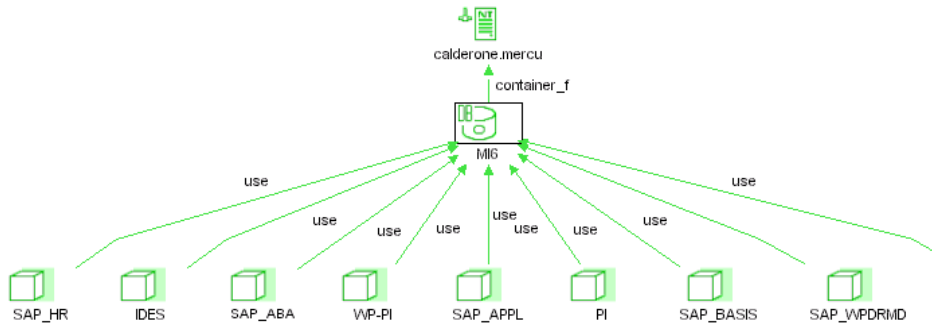
➤ **Symmetric Layout –** Enables you to display clear representations of complex networks:
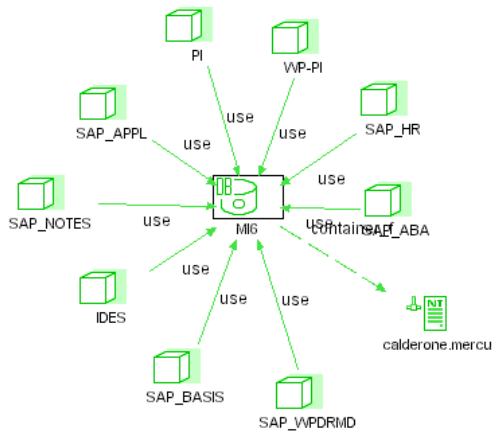
➤ **Manual Layout –** Enables you to manually change the topology map display, by dragging the displayed nodes and relationships to your required location.
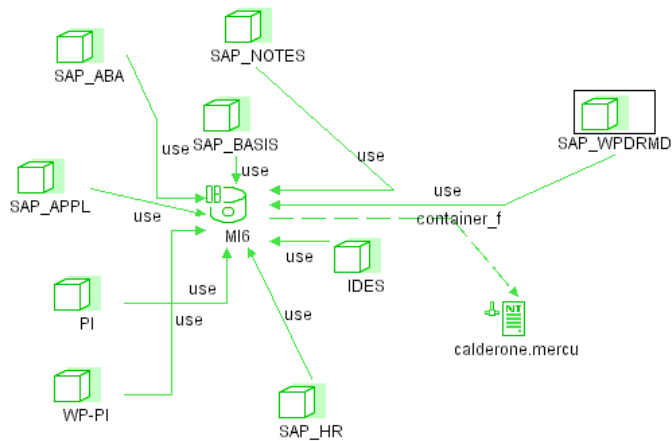
➤ **Orthogonal Layout** – Enables you to display views of outstanding clarity that are achieved by employing only horizontal and vertical edge routing.
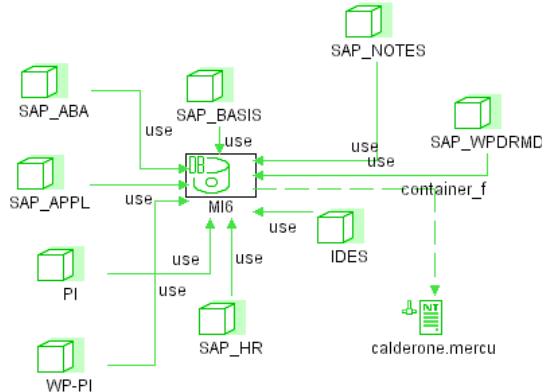
Note: You can save the manual changes you made, by clicking the **Save** button on the toolbar. Thus, you are able to reuse the layout you created.

You can only save one manual layout. If another manual layout is already saved, after you click the **Save** button a message is displayed asking whether you want to replace the saved layout. Click **Yes** to save the new layout. Click **No** to keep the old layout.

➤ **Zoom** – Zooms in on a specific section of the topology map by clicking and drawing a selection rectangle around the part you want to magnify. When you select this option, the pointer becomes a zoom.

➤ **Interactive Zoom –** When you select this option, the pointer becomes an interactive zoom. You decrease the view magnification by clicking and pushing the pointer away from you and magnify the view by clicking and pulling the pointer towards you.

➤ **Fit to Window –** Enables you to fit the complete view area in the window.

➤ **Information Browsers** – Enables you to open or close the following Information tabs, which are displayed in the Information pane in the topology view:

  ➤ "Configuration Item Browser" on page 167

  ➤ "Configuration Item Events" on page 239

  ➤ "View Events" on page 239

  ➤ "Inbox Tab" on page 239

➤ **Show Hidden Symbols** – Enables you to display hidden CIs.

➤ **Details –** Enables you to display the attributes of all CIs that are located in the same view layer in one table. Data is organized by the CIT of the CI. For details, see "Viewing Layers with a Large Number of CIs" on page 203.

Note: The **Zoom**, **Interactive Zoom**, **Fit to Window**, and **Layout** options are also displayed in a menu when right-clicking inside any empty area of the topology map.

Configuration Management Menu

The **Configuration Management** menu contains the following options:

➤ **Compare Compound** – Compares two compound objects, enabling you to see the difference between the hierarchies of the two objects. For details, see Chapter 25, "Comparison Capabilities"

➤ **System Reports** – Displays the System Reports dialog box, enabling you to select the report you want to view. For details, see "Viewing and Saving the Generated Report" on page 357.

➤ **Gold Master Reports –** Enables you to generate a report that compares the configuration of a GoldMaster CI to other CIs of the same CIT (for details, see "Gold Master Comparison" on page 423.

➤ **Impact Analysis** – Displays the Ticket Manager which enables you to simulate how changes you want to make will impact your system. For details on impact analysis, see "Evaluating the Impact of Infrastructure Events".

Administration Menu

The **Administration** menu contains the following options:

➤ **Discovery Manager** – Displays the Discovery Management dialog box, enabling you to design, edit, and activate discovery patterns. For details, see "Discovery Manager" in the *Mercury Application Mapping Administration Guide*.

➤ **User Management**

   ➤ **User Profile** – Displays the User Profile dialog box, enabling you to display in the topology map selected views and the events related to those views, use existing Event Filters to filter the displayed events, and define your display definitions. For details, see the *Mercury Application Mapping Administration Guide*.

   ➤ **User Manager** – Displays the User Manager dialog box, enabling you to create and manage the system's users.

   ➤ **Role Manager** – Displays the Roles dialog box, enabling you to create and manage roles. For details on users and roles, see Chapter 1, "User and Role Management," in the *Mercury Application Mapping Administration Guide*.

   ➤ **Security Manager** – Displays the Security Manager dialog box, enabling you to define specific user or role permissions for different Mercury Application Mapping resources. For details, refer to Chapter 1, "User and Role Management," in the *Mercury Application Mapping Administration Guide*.

   ➤ **Show Concurrent Users** – Displays the Show Concurrent Users dialog box, enabling you to see which users are currently logged in.

➤ **Set Active Server** – Displays the Set Active Server dialog box, enabling you to define which of the connected servers is the active one.

When working in the topology map or in other Managers, except for the CI Type Manager, you can work simultaneously with all servers to which you are connected. However, when using some of the menu's commands (**Map > Import**, **Export All**; **Edit > Filter Editor**, **Find**; **Administration** > all commands), or when working in the CI Type Manager, you can work only with the defined active server and the actions you perform relate only to it.

➤ **State Manager**– Displays the State Manager dialog box, enabling you to define a CI's management state. For details, refer to "Defining Attribute States" in the *Mercury Application Mapping Administration Guide*.

➤ **System Type Manager**– Displays the Enumeration Manager dialog box, enabling you to define an attribute whose value has been predefined. For details on creating Enumeration definitions, see "Enumerations and Lists" on page 67.

➤ **Event Configuration** – Displays the Event Configuration window, enabling administrators to configure the event system, which consists of defining rule and time actions, as well as the message catalog and priorities.

Note: User management and event configuration can be performed only by system administrators.

➤ **Scheduler** – Displays the Scheduled Action Manager dialog box, enabling you to schedule the system to automatically activate actions on a periodic basis. For example, you can define a schedule for an automatic backup of the database, and determine how many backups are to be saved. For details on scheduled actions, see "Task Scheduling" in the *Mercury Application Mapping Administration Guide*.

Help Menu

The **Help** menu contains the following option:
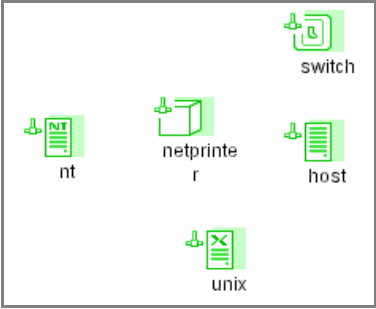
➤ **Help Topic** – Opens the Documentation Index.

➤ **About Mercury Application Mapping** – Displays software version information.

Toolbar

The toolbar buttons provide quick access to the most commonly used functions in the Mercury Application Mapping main window. The following table contains a brief description of each toolbar option.

Option Name	Button	Use This Option to...
New		Define the following: ➤ New items in all Mercury Application Mapping Managers (apart from the topology map) ➤ New queries in the TQL Builder. For details, see "Creating a TQL Query" on page 59. ➤ New views in the Service View Manager. For details, see "Creating a Pattern View" on page 110. ➤ New correlations rules in the Correlation Manager. For details, see "Creating a Correlation Rule" on page 284 ➤ New enrichment nodes and relationships in the Enrichment Manager. For details, see "Creating Enrichment Nodes and Relationships" on page 310. ➤ New System Reports in the Report Manager. For details, see "Creating a Report" on page 349. ➤ New CITs in the CI Type Manager. For details, see "Creating CITs" on page 384.
Open		Open a selected parent CI and drill down to its lower map layers in the topology map. For details, see "Navigating in Maps" on page 178.
Up		Move up one level in the Topology View. For details, see Part IV, "Viewing Queries with Topology View."
Back		Return to the last view you displayed.
Forward		Display a view you displayed before clicking the **Back** button.

Option Name	Button	Use This Option to...
Save		Save new and changed items in all Mercury Application Mapping Managers (apart from the CI Type Manager):
		➤ TQL queries in the TQL Builder. For details, see "Saving a Query" on page 90.
		➤ Managed views in the Service View Manager. For details, see "Saving the Pattern View" on page 159.
		➤ Correlations rules in the Correlation Manager. For details, see "Saving the Correlation Rule" on page 291.
		➤ Any changes to the manual layout in the topology map. For details, see "Improving Views" on page 197.
		➤ System Reports in the Report Manager. For details, see "Saving the System Report Template" on page 356.
Print		Print the contents of the topology map and the contents of the Information pane in the topology map. For details, see "Printing the Topology Map" on page 220.
Delete		Delete a selected item such as a view or TQL from the TQL Builder or Service View Manager or a CIT from the CI Type Manager.
		When you delete a CIT from the CI Type Manager, the CIT no longer appears in the CI Type Model, the Service View Manager, the Enrichment Manager, and the TQL Builder.
		To restore it, you must redeploy the package that contains the deleted CIT. For details, see "Deploying Packages With the JMX Console" in the *Mercury Application Mapping Administration Guide*.
Layout	Hierarchical Layout	Select the layout display from the list. For further details about each of the options, see "View Menu" on page 28.
Reorganize layer		Reorganize the CIs/CITs/nodes according to the selected layer. This is used when CIs/CITs/nodes have been moved around.
Refresh		Reorganize the CIs according to the selected layer. This is used when CIs have been moved around or when hidden CIs exist (used mainly for the topology map).

Option Name	Button	Use This Option to...
Rebuild		Rebuild the view in the Topology View manager (used when the view name is selected in the View Explorer).
Show Hidden Symbols		Display previously hidden CIs.
Select All		Select all the CIs and relationships in a selected layer.
Select All Nodes		Select all the nodes in a selected layer.
Hide All Relationships		Toggle between showing or hiding the relationships in the displayed view.
Hide All Relationship Labels		Toggle between showing or hiding the relationship labels in the displayed view.
Hide All Node Labels		Toggle between showing or hiding the node labels in the displayed view.
Set Nodes Image Size		Toggle between displaying a map with large icons with labels and displaying a map with small icons without labels (the labels are hidden because they become too small to read). Useful to display large maps..

large icons

small icons

Option Name	Button	Use This Option to...
Zoom		Zoom in on a specific section of the topology map by clicking and drawing a selection rectangle around the part you want to magnify. The area you have selected is displayed at the highest percentage that fits in the topology map.
Interactive Zoom		➤ Decrease the view magnification by clicking and pushing the pointer away from you. ➤ Magnify the view by clicking and pulling the pointer towards you.
Pan		Drag the view.
Navigation		Take the pointer to the next connected node in a clockwise direction. ➤ Left-click to go from node to node in a clockwise direction. ➤ Right-click to point to the next node in a clockwise direction. **Note:** You can use this option only on nodes that are connected by relationships.
Select Mode		Select a CI or CIs.
Fit to Window		Fit the complete view in the window.
Place in Center		Places the selected CI in the center of the map.
Create Relationship		Create a relationship between two existing CIs. For details, see"Adding Nodes and Relationships to a TQL Query" on page 116.
TQL result count		Calculate the number of instances found for each CI/TQL node/relationship.

Option Name	Button	Use This Option to...
Information Pane		Display or hide the Information pane (containing the **Configuration Item Browser**, **Configuration Item Events**, and **View Events** tabs in the topology map).
Explorer Pane		Display or hide the View Explorer in the selected tab.
Filters Pane		Display or hide the Tags pane in the selected tab.
Map Pane		Display or hide the topology map in the selected tab.
Help Topic		Display the table of contents for the online Help.

The Main Window's Level Structure

Mercury Application Mapping Main Window includes the Folder, View, and CI levels. They can include additional sublevels or layers.

To drill down from a higher level to a lower one, double-click the symbol on the topology map. For more navigation options, see "Navigating in Maps" on page 178.

This section includes:

➤ "Folder Level" on page 39

➤ "View Level" on page 40

➤ "CI Level" on page 41

Folder Level

If you choose to connect to only one server in the Login stage, or if you double-clicked one of the displayed servers, you can view either the Folder level or the View level (see next section), depending on whether views are organized in folders in your system:
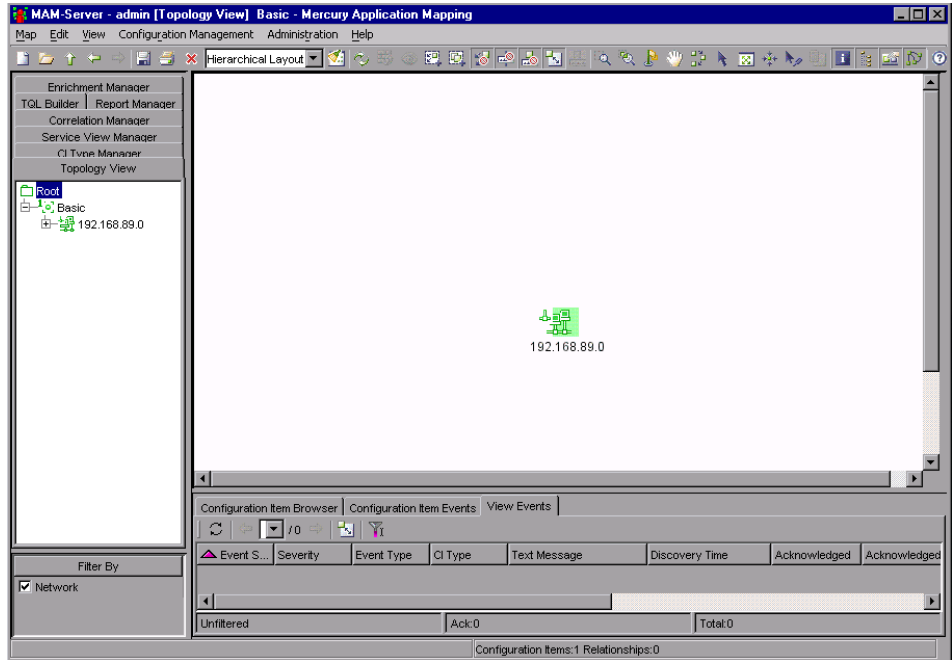
A folder can contain subfolders and several views. You use the Service View Manager to organize the views in a folder. For details, see "View Shortcut Menu" on page 103.

View Level

The View level displays the views that are set in the Service View Manager to be displayed on the topology map, and the views that are part of your user profile. For details, see "Understanding Layer Layouts" on page 179.
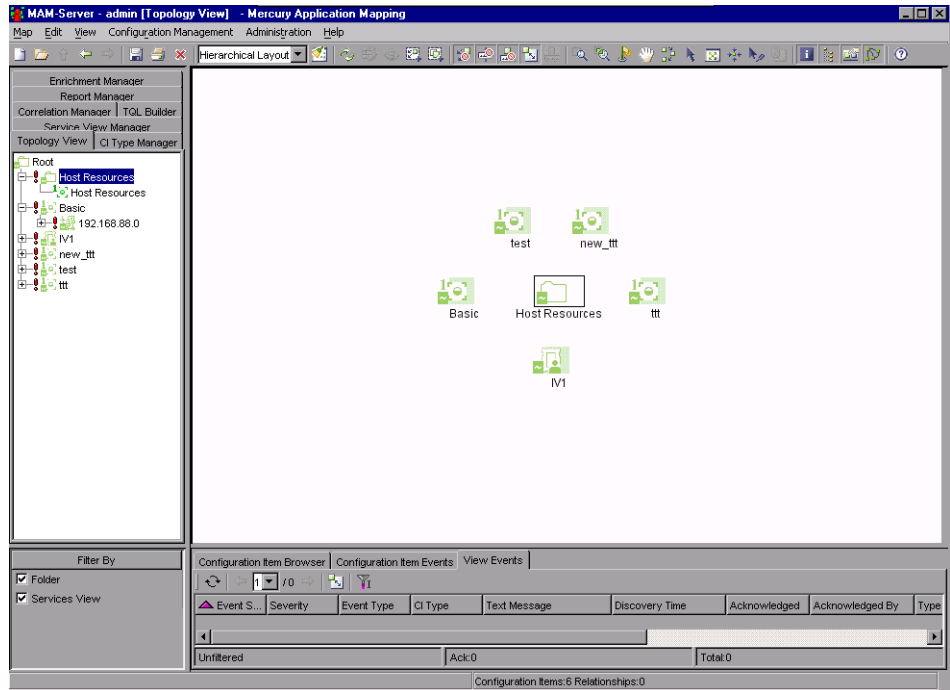
Note: A view's icon can be customized in the Service View Manager. For details, see "Navigating in Maps" on page 178.

CI Level

The CI level displays the view results, that is, the CIs and relationships that fulfill a View TQL query. These CIs and relationships are displayed according to display and organizational rules that are assigned to them in the Service View Manager:
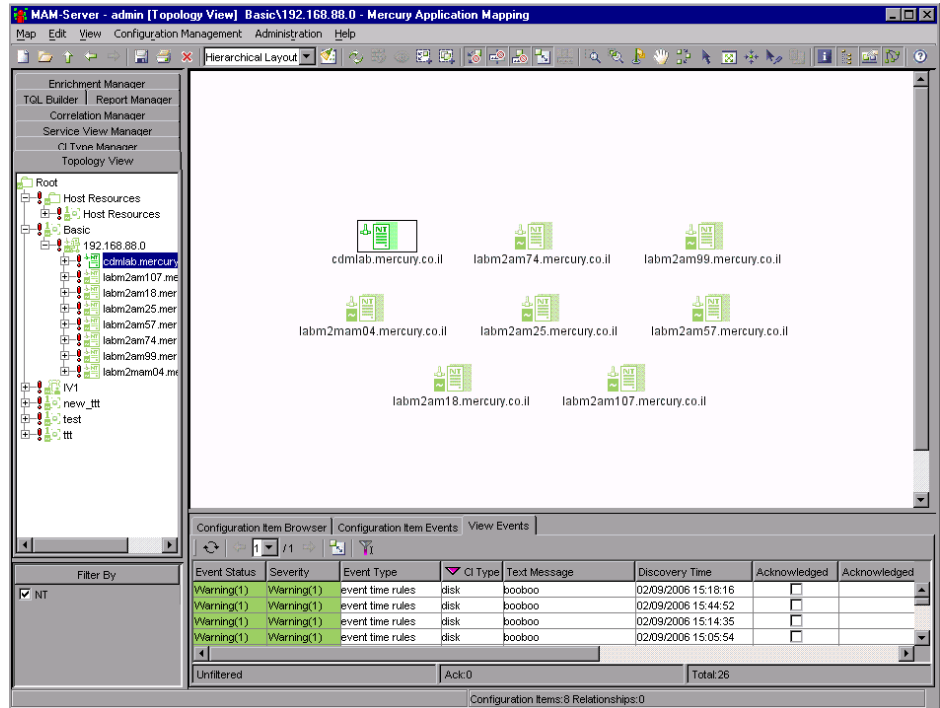
The icon that represents each CI and relationship is determined by the CI/relationship CIT. CI icons are made up of the following three components:

➤ An icon.

➤ The group to which it belongs for example, Network or SAP, appears in the upper left-hand corner.

➤ The severity for example, critical or major, appears in the lower left-hand corner.

If an icon appears with a shadow, as shown in the example above, it means that there is an additional CI layer beneath it. Double-click to drill down to the lower level(s).

Working with Mercury Application Mapping

Several features of Mercury Application Mapping work the same way across the different components:

➤ "Selecting Multiple CIs" on page 43

➤ "Working with Shortcut Menus" on page 44

➤ "Tooltips" on page 44

➤ "Drilling Down" on page 44

➤ "Operator List" on page 44

Selecting Multiple CIs

In the topology map you can select more than one CI.

To select multiple CIs:

➤ Hold down CTRL and click the CIs or drag the mouse diagonally along the CIs as shown in the following figure:

Before   After
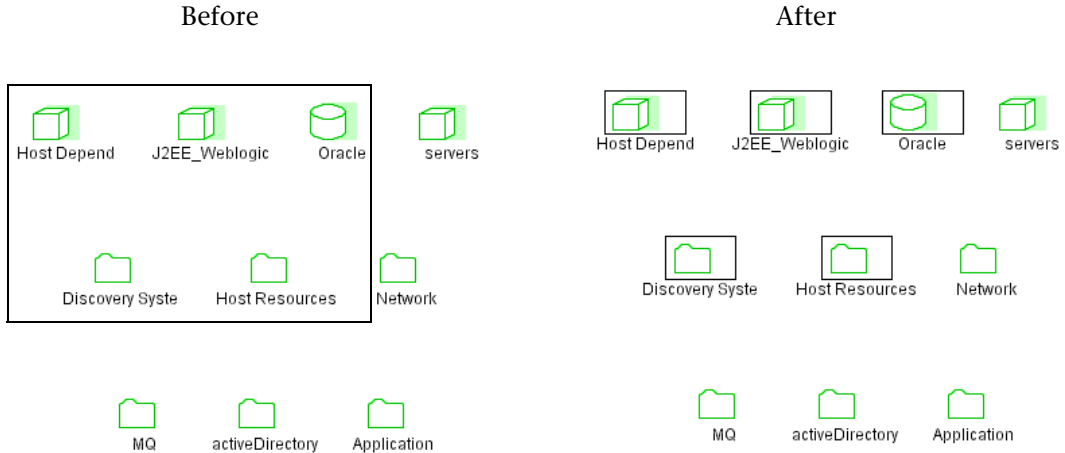
Working with Shortcut Menus

Right-clicking a CI in the topology map opens a shortcut menu. The shortcut menu depends on the selected CI. The shortcut menus are detailed in each relevant part of this guide.

Tooltips

When the pointer is moved over a CI in the topology map, a tooltip displays information that is relevant to the CI.

Drilling Down

If a CI icon appears with a shadow, double-click it to display the layer beneath it.

Operator List

The operators in the **Operation** list are:

Operator	Description
Equal	Checks whether the attribute value is equal to the value specified in the value box.
Equal ignore case	Checks whether the attribute value is equal to the value specified in the value box independently of the case.
Not Equal	Checks whether the attribute value is not equal to the value specified in the value box.
Greater	Checks whether the attribute value is greater than the value specified in the value box.
Greater than or equal	Checks whether the attribute value is greater than or equal to the value specified in the value box.
Less	Checks whether the attribute value is less than the value specified in the value box.
Less than or equal	Checks whether the attribute value is less than or equal to the value specified in the value box.
Like	Uses a wildcard (%). Use Like when you are not sure of the complete name of what you are looking for.

Operator	Description
Like ignore case	Uses a wildcard (%). Use Like ignore case when you are not sure of the complete name of what you are looking for. The case of the string is ignored.
In	Displays only the instances where this attribute value equals one of the selected values. For example, for CIs that have an operational state that equals Warning(2) and Warning(3), select the operator **In**, and select both Warning(2) and Warning(3) from the **Value** list.
Is null	Checks whether the attribute value is null.
Change during	(displayed only when you select the **Create Time** attribute) Displays only the instances that changed during the period specified in the Value box.
Unchanged during	(displayed only when you select the **Create Time** attribute) Displays only the instances that did not change during the period specified in the Value box.

Part II

Topology Query Language Builder

3

Introducing the Topology Query Language (TQL) Builder

This chapter introduces the Topology Query Language (TQL) Builder, which enables you to query the CMDB by defining TQL queries.

This chapter describes:	On page:
About the Topology Query Language (TQL) Builder	50
Working With TQL Builder	51
TQL Builder Shortcut Menus and Toolbar Options	52

About the Topology Query Language (TQL) Builder

The TQL Builder enables you to define and save different types of TQL queries. These queries, which use a structured language similar to SQL, perform various functions in the system, such as extracting the information from the CMDB that is relevant to the user, or creating a report containing specific information. Once a TQL query has been created, it resides in the system memory and generates updated results automatically. For details, see"Mercury Application Mapping Topology Query System" on page 9.
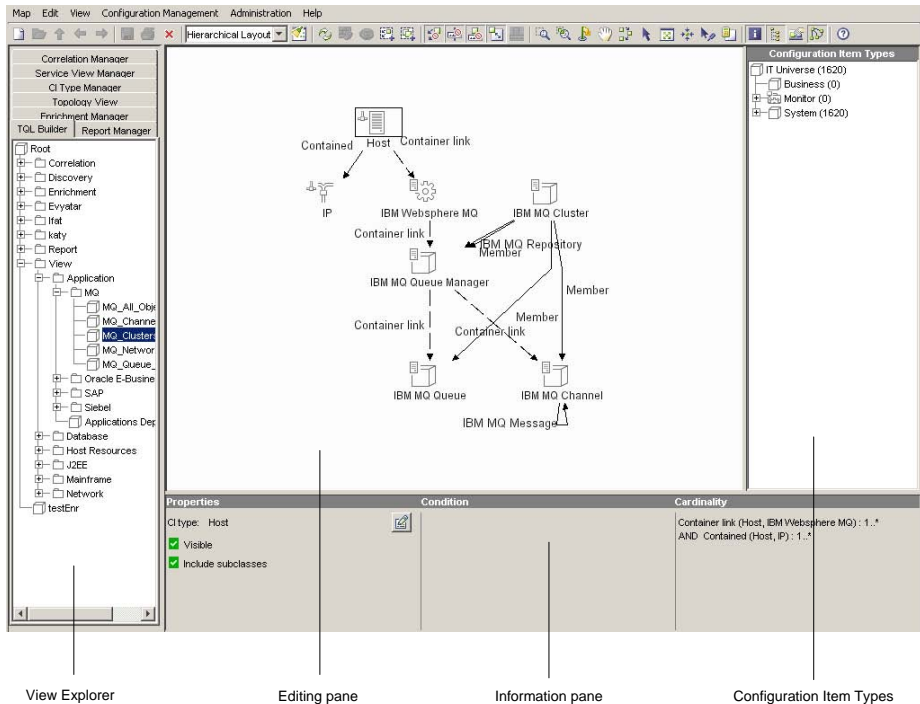
| View Explorer | Editing pane | Information pane | Configuration Item Types |

When the TQL Builder tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **View Explorer** – Displays a hierarchical tree structure of the queries that you defined. You can copy, export, and save TQL queries as XML scripts to be used for backup purposes.

Note: Each TQL definition type is represented by a unique icon.

➤ **Editing pane –** Displays the currently selected TQL query, which consists of TQL nodes and the relationships between them.

➤ **Configuration Item Type –** Represents the CI Type Model and contains icons for each CIT, as defined by the administrator (for details, see "Assigning an Icon to a CIT" on page 389). By clicking and dragging nodes to the editing pane and then defining the relationship between them, you can define a query and save it to the database. For example, you can drag the NT and IP nodes to the editing pane and then define the connection between them by adding relationships.

➤ **Information Pane –** Displays the attribute conditions you defined for the selected node and relationship. For details, see, "Setting TQL Node and Relationship Definitions" on page 68.

Working With TQL Builder

Use the TQL Builder to perform the following functions:

➤ Create various types of TQL queries, as well as duplicate and delete TQL queries.

➤ Delete selected TQL nodes and relationships in the editing pane.

➤ Define the relationship between TQL nodes by adding relationships, either manually (one by one) or by selecting multiple relationships to speed up the process.

➤ Define the specific attribute conditions of TQL nodes and relationships in the query.

➤ Define the attributes whose information should automatically be updated in the topology map whenever the attribute is changed.

➤ Define whether to display the result of the query in the topology map, including the CIs derived from the selected TQL nodes.

➤ Export and import selected queries to and from XML script format files, for example, to relocate TQL queries from one workstation to another.

➤ Zoom in and out of the editing pane to view the selected TQL at different zoom levels.

➤ Display the queries in different layout views.

TQL Builder Shortcut Menus and Toolbar Options

The TQL Builder contains different right-click shortcut menus, depending on your selection, as follows:

➤ "TQL Toolbar Options" on page 52"

➤ "TQL Folder Shortcut Menu" on page 53"

➤ "TQL Query Menu" on page 53

➤ "Multiple TQL Node Menu" on page 54

➤ "Single TQL Node Shortcut Menu" on page 55

➤ "Relationship Shortcut Menu" on page 56

TQL Toolbar Options

For a description of each toolbar option in the TQL Builder, see "Toolbar" on page 34.

TQL Folder Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a TQL folder in the View Explorer:

Option	Description
New	Enables you to define a new TQL Query.
New Folder	Enables you to create a new folder.
Delete	Enables you to delete a selected TQL folder.

TQL Query Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a TQL query in the View Explorer:

Option	Description
New	Enables you to define a new TQL Query. For details, see "Creating a TQL Query" on page 59.
New Folder	Enables you to create a new TQL folder.
Save	(Enabled only when a new TQL is created or when changes are made to an existing TQL.) Enables you to save the selected query to the database, as described in "Saving a Query" on page 90.
Save As	Displays the Save As dialog box, enabling you to define a new name and description for a query. This option can be used for copying queries, as described on "Creating a Query by Copying an Existing Query" on page 60.
Delete	Enables you to delete a selected TQL query.
Export	Displays a standard Save As dialog box, enabling you to save the TQL query in XML script format. This option can be used to move TQL queries from one workstation to another.
Properties	Enables you to change the description and priority of the selected TQL query, but not the name of the query or its type.

Tip: Hold the pointer over a TQL node or relationship to view a tooltip. The tooltips for the TQL nodes and relationships contain the definitions of the selected TQL nodes, for example, the attribute conditions, as described in "Setting TQL Node and Relationship Definitions" on page 68.

Multiple TQL Node Menu

The following table provides a brief description of each option in the shortcut menu displayed by selecting two or more TQL nodes in the editing pane. For a description of how to select multiple TQL nodes, see "Selecting Multiple CIs" on page 43.

Option	Description
Add Relationship	(For two selected TQL nodes only) Displays the Add Relationship dialog box, enabling you to add a predefined relationship between two TQL nodes in a query, as described in "Adding Nodes and Relationships to a Query" on page 62. (Only relevant relationships are displayed in the **Add Relationship** list.)
Delete	Enables you to delete the selected TQL nodes.
Copy	Enables you to copy the selected TQL nodes.

Single TQL Node Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a single TQL node in the editing pane:

Option	Description
TQL Node Wizard	Builds a TQL query. For details, see "Using the TQL Node Wizard" on page 80.
Add Relationship	Displays the Add Relationship dialog box, enabling you to create a relationship for your TQL nodes by selecting it from a predefined list. For details, see "Adding Nodes and Relationships to a Query" on page 62.
TQL Node Definition	Displays the TQL Node Definition dialog box, enabling you to define the attribute conditions for the selected TQL node. For details, see "Setting TQL Node and Relationship Definitions" on page 68.
Delete	Enables you to delete the selected TQL node(s).
Add SubGraph	Creates a graph that represents additional TQL query data related to a specific node. For details, see "Creating a Dependency Graph" on page 86.
Show element's instances	Displays all of the instances found for each TQL node in a table. For details, see "Showing TQL Node Instances" on page 88.
Copy	Enables you to copy the selected TQL nodes.

Relationship Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a relationship in the editing pane:

Option	Description
TQL Relationship Definition	Displays the TQL Relationship Definition dialog box, enabling you to define the attribute conditions for the selected relationship. For details, see "Setting TQL Node and Relationship Definitions" on page 68.
Delete	Enables you to delete the selected relationship.
Straighten	Enables you to straighten the relationship between two nodes. This option is only available for relationships that have angles. For details, see "Straightening a Relationship With an Angle" on page 86.
Copy	Enables you to copy the selected TQL node(s).

4

Defining Topology Query Language (TQL) Queries

This chapter explains how to define TQL queries.

About Defining TQL Queries

When defining a new query with the TQL Builder, you first select the type of query to create, and then select the TQL nodes that are to be part of the query. You then define specific attribute conditions for each node, including the attributes that define the relationships between nodes.

You can either define the TQL by manually dragging and dropping the TQL nodes manually from the Configuration Item Types pane or use the TQL Node Wizard (for details, see "Using the TQL Node Wizard" on page 80).

Note: A TQL query is subject to certain validation restrictions. For details, see "Validation Restrictions" on page 93.

TQL Query Workflow

You create TQL queries according to the following workflow:

➤ Create a new TQL query.

For details, see "Creating a TQL Query" on page 59.

➤ Add nodes and relationships to the query.

For details, see "Adding TQL Nodes and Relationships to a Query" on page 61.

➤ Define node and relationship attribute conditions.

For details, see "Setting TQL Node and Relationship Definitions" on page 68.

➤ Create a dependency graph.

For details, see "Creating a Dependency Graph" on page 86.

➤ Save the new query.

For details, see "Saving a Query" on page 90.

Creating a TQL Query

You can create a new query or you can use an existing TQL, from any query type, as a basis for your new TQL. In either case you must give a unique name to the query.

This section includes:

➤ "Creating a New Query" on page 59

➤ "Creating a Query by Copying an Existing Query" on page 60

Creating a New Query

You can create a new query.

To create a new query:

 1 Click **Map** > **New** button on the toolbar or right-click the folder in which you want to create the new query, and select **New** to open the Create New TQL dialog box.

Note: You can, at a later stage, change the query's location by dragging and dropping the folder into another folder.

---

 2 In the **TQL Name** box, type a unique name for the TQL query.

 3 From the **Type** list, choose the type of TQL query you want to create, as follows:

➤ **View** – Creates queries whose results are displayed in the Topology View manager. For details, see Part IV, "Viewing Queries with Topology View."

➤ **Correlation** – Creates queries regarding correlation nodes that influence other correlation nodes in the system, as defined using the Correlation Manager. For details, see Part VI, "Correlation Rules."

➤ **Enrichment** – Create queries regarding enrichment nodes and relationships, whose existence can be deduced from prior or accumulated knowledge about your system. For details, see Chapter 17, "Defining Enrichment Nodes and Rules."

➤ **Report** – Creates queries whose results are compiled into reports using the Report Manager. For details, see Chapter 19, "Introduction to the Report Manager."

➤ **Discovery** – Creates queries for use by the discovery system. For details, see, "Discovery Manager" in the *Mercury Application Mapping Administration Guide*.

Note:

➤ Discovery TQLs should only be created by users who are thoroughly familiar with the discovery patterns used by the discovery system. For details, contact Mercury Customer Support.

➤ For the Correlation, Enrichment and Report TQLs to be valid, they must comply with the restrictions described in "Validation Restrictions" on page 93. If the TQL is not valid, you cannot save it.

---

 4 From the **Priority** list, select a priority level for the new TQL query (**Low, Medium**, **High** or **Express**). This setting determines how often the query should be rerun automatically by the system to include updated information from the database.

 5 If necessary, in the **Description** box, enter a description of the TQL query.

 6 If necessary, select **Persistent** to define whether the selected TQL should always remain in the system memory. You should use this option for TQLs that are used frequently, and not for TQLs that are only used occasionally.

 7 Click **OK** to open the new TQL query in the View Explorer.

Creating a Query by Copying an Existing Query

You can create a new query by copying an existing query and modifying it.

To create a query by copying an existing one:

 1 On the View Explorer, right-click the query you want to use as a basis for the new query. Click **Save As** to display the Save TQL as dialog box.

 2 In the **TQL Name** box, enter a unique name for the new TQL query.

3 From the **Type** list, select the type of the new TQL query. For details, see step 3 in "Creating a New Query" on page 59.

4 Follow steps 4 to 7 in "Creating a New Query" on page 59.

TQL Builder Tooltips

Hold the pointer over a node or relationship to view its tooltip. The tooltips contain the definitions of the selected nodes and relationships. For details, see for example, the attribute conditions, as described in "Defining TQL Nodes and Relationships" on page 69.

Adding TQL Nodes and Relationships to a Query

After you have created the TQL query, the next step is to add the TQL nodes and relationships that define the query. The TQL nodes represent the CITs, as defined in the CI Type Manager, and the relationships represent the connection between them. Relationships are defined one at a time for each pair of TQL nodes in the query.

For a list of the relationships you can use to link two nodes and their definitions, see "Relationship Definitions" on page 543.

This section contains the following topics:

➤ Adding Nodes and Relationships to a Query – for details, see page 62

➤ Defining Function and Compound Relationships – for details, see page 63

➤ Sample Compound Relationship – for details, see page 66

➤ Deleting a Node or Relationship – for details, see page 68

➤ Editing a Relationship Type – for details, see page 68

Adding Nodes and Relationships to a Query

This section describes how to add TQL nodes and relationships to a query.

To add TQL nodes and relationships to a query:

1 From the tree in the View Explorer, select the TQL query to which you want to add TQL nodes and relationships.

2 From the tree displayed in the Configuration Item Type pane, click and drag one or more required TQL nodes on to the Editing pane. These are the TQL nodes that are included in the query.

3 To add a relationship between two nodes:

➤ Select the required TQL node(s) by holding down CTRL and clicking the TQL nodes, right-click and select **Add Relationship**.

or

➤ Click the **Create Relationship** icon and draw a line between the required nodes.

The Add Relationship dialog box contains the following fields:

Field	Description
Node Label	The label of the selected nodes.
Relationship Direction	The direction of the relationship that indicates which node is dependent on the other.
Relationship Type	A valid relationship that defines the connection between the selected nodes.
Advanced	Define the connection between two nodes using either a function or compound relationship
	You can define the connection between nodes by adding either function or compound relationships. For details, see "Defining Function and Compound Relationships" on page 63.

4 Select the required direction of the relationship.

5 If necessary, click **Advanced** to select either function or compound relationships.

6 Click **OK**. The selected nodes are linked by the relationship you have selected.

The direction of the relationship indicates which node is dependent on the other. The following example displays two hosts (an NT and a Unix), that have a client and a server port, and are linked to one another via a Client-Server connection.

Defining Function and Compound Relationships

You can define the connection between two nodes using either a function or compound relationship.

➤ A function relationship (**Virtual join link**) describes a logical connection between two nodes, which is not represented in the CMDB.

When you create a function relationship, you must define an attribute for each node, whose value will be used for comparison.

➤ A compound relationship (**Virtual compound link**) describes a connection that is composed of more than one relationship. If there are two nodes whose connection cannot be directly defined by one relationship, you can add several relationships, which act as stepping stones along the line, to create a connection.

To define function and compound relationships:

1 Select the nodes you want to connect with a function or compound relationship (for details on how to add nodes to a TQL query, see "Adding TQL Nodes and Relationships to a Query" on page 61).

2 Right-click and select **Add Relationship** to open the Add Relationship dialog box.

3 Click **Advanced**.

4 From the **Relationship** list, select **Function Relationship**.

5 To add a **Virtual join link**, do the following:

➤ Select **Virtual - Join link**.

➤ In the **Join Definition** area, click the **Adds an attribute definition** button to open the Conditions dialog box.

➤ From the **Attribute** lists, choose the two required attributes.

➤ From the **Operator** list, select one of the following operators:

• Select **Equal** if you want the system to check whether the two selected attributes are equal.

• Select **Not equal** if you want the system to check whether the two selected attributes are not equal.

➤ Click **OK** to save your definition. The definition appears in the **Join Definition** area.

➤ To delete a definition, select the required definition and click the **Deletes selected row** button.

➤ To edit a definition, select the required definition and click the **Edits selected row** button to open the Conditions dialog box. Edit the definition as described in step 5.

➤ Click **OK** to save your changes.

6 To add a **Virtual compound link**, do the following:

➤ Select **Virtual - Compound link.**

➤ In the **Compound Definition** area, click the **Add** button to open the Add dependency dialog box.

➤ From the **Source** list, select the source node.

➤ From the **Target** list, select the target node.

➤ From the **Relationship** list, select an available relationship connecting the two nodes. The list of available relationships appears only after defining both nodes.

➤ From the **Relationship Direction** list, select the required direction.

Note: You can create as many virtual compound definitions as necessary.

➤ Click **OK** to save the compound definition. The compound definition appears in the Compound definition dialog box.

Note: Each node represents a step in the chain leading from one node to the other. Ensure that **<node 1>** of the first row and **<node 2>** of the last row are the two selected nodes in the TQL query.

➤ To delete a definition, select the required definition and click the **Delete** button.

➤ To edit a definition, select the required definition and click the **Edit** button to open the Add dependency dialog box. Edit the definition as described in step 6.

➤ In the **Depth** box, select a number that determines how many steps (that is, how many connecting nodes) in the TQL pattern you want included in the discovery process. The default is 5.

7 Click **OK** to save your definitions.

Sample Compound Relationship

This section contains a sample compound relationship comprising several nodes and their connections, displayed in two different views:

➤ A normal view in which each node in the compound relationship was dragged and dropped from the Configuration Item Types pane. For details, see "Adding TQL Nodes and Relationships to a Query" on page 61.

The normal view result is displayed below:

➤ A compound view in which the compound relationship was defined by using the **Compound Definition** area of the Add Relationship dialog box, as illustrated below. (For details, see "Defining Function and Compound Relationships" on page 63.)

The following compound view is created:

The compound view result is as follows:

Deleting a Node or Relationship

This section describes how to delete a node or relationship.

To delete a node or relationship:

Right-click the TQL node or relationship you want to delete and select **Delete**.

Editing a Relationship Type

This section describes how to edit an existing relationship. This option only appears if at least one of the following is true:

➤ You can change the direction the relationship.

➤ You can select a descendent of the existing relationship.

➤ The selected relationship is a function or compound relationship.

To edit a relationship type:

1 Right-click the required relationship and select **Edit Relationship Type**. The Edit Relationship Type dialog box opens.

2 Make the required changes. For details, see "Adding TQL Nodes and Relationships to a Query" on page 61 and "Defining Function and Compound Relationships" on page 63.

3 Click **OK** to save your changes.

Setting TQL Node and Relationship Definitions

After you have added the TQL nodes and relationships required for your query, you can define their specific attribute conditions.

This section includes:

➤ Defining TQL Nodes and Relationships – for details, see page 69

➤ Changing the CI Type of the TQL Node – for details, see page 70

➤ Defining Attribute Conditions for Nodes and Relationships – for details, see page 71

➤ Filtering Query Results in the Attribute Condition Tab – for details, see page 73

➤ Defining Qualifier Conditions – for details, see page 75

➤ Defining Relationship Cardinality – for details, see page 76

➤ Filtering Query Results in the Cardinality Tab – for details, see page 77

Defining TQL Nodes and Relationships

This section describes how to define TQL nodes and relationships.

To define TQL nodes and relationships:

1 In the Editing pane, right-click the TQL node or relationship whose attribute conditions you wish to define, and select **TQL Node Definition** or **TQL Relationship Definition** to open the TQL Node/Relationship Definition dialog box.

Note: Alternatively, you can select the required node/relationship, and click the **Go to Edit Dialog** button in the Information pane to open the TQL Node/Relationship Definition dialog box.

---

2 (Optional) The **Element name** box contains the name of the selected node/relationship. By default, Mercury Application Mapping assigns the CIT type as the element's name. You can rename a TQL node in the **Element name** box giving it a unique label. This can be helpful when there is more than one node of the same CIT type in the TQL query.

3 To display the selected TQL node in the Editing pane, select **Visible**. When **Visible** is cleared, an invisible box appears to the right of the selected TQL node in the editing pane:

Any query results pertaining to that TQL node are not displayed in the topology map. This can be useful when certain relationships or TQL nodes are required to build the query but are not needed in the results. For example, NTs are connected to the network by defining specific IPs, but in any query results you may want to view the IP elements only and not the NT elements

4 To display both the selected CI and its descendents in the topology map, select **Include subclasses**.

Changing the CI Type of the TQL Node

You can change the CI type of the TQL node after you have created a TQL.

To change the CI type of a TQL node after creating a TQL query:

1 In the Editing pane, right-click the TQL node whose CI type you want to change. and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Change CI type** button to open the Change CI type dialog box.

3 Select the required CI type and click **OK**.

Note: You can only change the CI type to a type from one of the CIT's subclasses, if one exists.

Defining Attribute Conditions for Nodes and Relationships

This section describes how to define attribute conditions for nodes and relationships.

To define attribute conditions:

1 In the Editing pane, right-click the TQL node or relationship whose attribute conditions you wish to define, and select **TQL Node Definition** or **TQL Relationship Definition** to open the TQL Node/Relationship Definition dialog box.

2 In the TQL Node/Relationship Definition dialog box, click the **Attribute Conditions** tab.

3 Click the **Add an attribute condition** button to open the Condition dialog box.

4 From the **Attribute Name** list, choose the required attribute.

5 From the **Operation** list, select the required operation.

The operations in the **Operation** list are:

Operation	Description
Equal	Checks whether the attribute value is equal to the value specified in the **Value Comparison** box.
Not equal	Checks whether the attribute value is not equal to the value specified in the **Value Comparison** box.
Is null	Checks whether the attribute value is null.
Like	Uses a wildcard (%). Use **Like** when you are not sure of the complete name of what you are looking for.
In	Displays only the instances where this attribute value equals one of the selected values. For example, for CIs that have an operational state that equals Warning(2) and Warning(3), select the operation **In**, and select both Warning(2) and Warning(3) from the **Value comparison** list.

Operation	Description
Equal ignore case	Checks whether the attribute value is equal to the value specified in the value box regardless of the case.
Greater	Checks whether the attribute value is greater than the value specified in the **Value comparison** box.
Great than or equal	Checks whether the attribute value is greater than or equal to the value specified in the **Value comparison** box.
Less	Checks whether the attribute value is less than the value specified in the **Value comparison** box.
Less than or equal	Checks whether the attribute value is less than or equal to the value specified in the **Value comparison** box.
Changed during	(Displayed only when you select the **Create Time** attribute.) Displays only the instances that changed during the period specified in the **Value comparison** box.
Unchanged during	(Displayed only when you select the **Create Time** attribute.) Displays only the instances that did not change during the period specified in the **Value comparison** box.
In list	Equals one of the elements that appears in the list. For example: 320,4445,3483.

Note: When using the **Not equal** or **Not like** operations, query results do not include data from the CI instances that were not assigned a value. For example, let us assume that your system contains three hosts. Host1 was assigned the value A, Host2 was assigned the value B, and Host3 was not assigned a value. If you created a query in which you want to retrieve all hosts that are **Not equal** to A, the results of your query only include Host2 because Host3 was not assigned a value.

6 If you want to make a value comparison, do the following:

➤ Select **Value**.

➤ In the **Value Comparison** box, enter or select the value of the attribute.

Note: If you use a string in the **Value** box, it must appear in single quotes, as follows: 'YOKO', 'JASON'.

---

7 If you want to create a condition in which you compare two attributes, do the following:

➤ Select **Attribute**.

➤ In the **Attribute Comparison** box, select the required attribute from the **Attribute Comparison** list.

8 To filter the query results, see "Filtering Query Results in the Attribute Condition Tab" on page 73.

9 Click **OK** to save the changes you have made and close the Condition dialog box.

Filtering Query Results in the Attribute Condition Tab

In the **Attribute Condition** tab, you can create an expression defining conditions that restrict the number of nodes that appear in the query.

To filter your query results:

1 In the Editing pane, right-click the TQL node or relationship whose attribute conditions you wish to define, and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Attribute Condition** tab.

3 In the **Condition** area, create an expression using the following:

Field	Description
NOT	Select **NOT** if you want the condition statement to do the opposite of what is defined.
	Note: If you select **NOT**, query results do not include data from the CI instances that were not assigned a value. For example, let us assume that your system contains three hosts. Host1 was assigned the value A, Host2 was assigned the value B and Host3 was not assigned a value. If you created a query in which you want to retrieve all hosts that are equal to A and selected **NOT** , the results of your query only include Host2 because Host3 was not assigned a value.
Brackets ⎡ { ⎤ ⎡ } ⎤	Click inside the **Brackets** box to display a list of brackets you can use to build more complex logical statements.
Criteria	Contains the attribute condition definition as defined in the Condition dialog box (for details, see step 3 in "Defining Attribute Conditions for Nodes and Relationships" on page 71).
And/Or	Click inside the **And/Or** field and select either **AND** or **OR** to link multiple conditions.

4 To delete an attribute condition, select the attribute condition you want to delete and click the **Delete selected row** button.

5 To insert an attribute condition before a selected row, select the required row, and click the **Insert an attribute condition before a selected row** button to open the Condition dialog box.

6 To edit an attribute condition, select the attribute condition you want to edit and click the **Edit selected row** button to open the Condition dialog box. For details on how to define an attribute condition, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

7 To move a row up, select the row you want to move and click the **Move selected row up** button.

8 To move a row down, select the row you want to move and click the **Move selected row down** button.

9 Click **OK** to save the changes you have made.

Defining Qualifier Conditions

This section describes how to define qualifier conditions for the selected node/relationship. For example, you can use a qualifier to define a CIT as abstract, meaning you cannot create instances from it.

To define qualifier conditions:

1 In the Editing pane, right-click the TQL node or relationship whose qualifier conditions you wish to define, and select **TQL Node Definition** or **TQL Relationship Definition** to open the TQL Node/Relationship Definition dialog box.

2 Click the **Qualifier condition** tab.

3 Select the required qualifier(s) according to the following:

➤ **NOTIFY_CHANGES** – For internal use only

➤ **BLE_LINK_CLASS** – For internal use only

➤ **ABSTRACT_CLASS** – You cannot create instances of this CIT

➤ **ITU_HIDDEN_CLASS** – For internal use only

➤ **CONTAINER** – Relevant for relationships only. Represents containment between two CIs.

➤ **HIDDEN_CLASS** – Does not appear anywhere in the application.

➤ **READ_ONLY_CLASS** – For internal use only

Note: You can hold the CTRL key down to make multiple selections.

➤ Click **OK** to save your definitions.

Defining Relationship Cardinality

In this section, you define how many CI instances you expect to have at the end of a relationship in your query result.

To define relationship cardinality:

1 In the editing pane, right-click the TQL node whose conditions you wish to define, and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Cardinality** tab.

Note: This tab only appears when a node is selected.

---

3 Click the **Add an attribute condition** button to open the Relationship Cardinality dialog box.

4 In the **Node** box, select the required relationship to which the selected node is attached. The list contains all the relationships that are linked to the selected node.

In the following steps, you define the lower and upper limits for including the node at the other end of the relationship in the query results.

5 In the **Min** box, enter the value that defines the lower limits for including the node at the other end of the relationship in the query results.

6 In the **Max** box, enter the value that defines the upper limits for including the node at the other end of the relationship in the query results. You can use an asterisk (*****) in the **Max** box to represent an infinite value.

For example, if <end_1> is IP and <end_2> is NT, entering 1 in the **Min** box and asterisk (*****) in the **Max** box instructs the system to retrieve only those IPs that are connected to at least one NT. (The asterisk indicates an infinite value.) Entering 3 in the **Min** box and asterisk (*****) in the **Max** box instructs the system to retrieve only those NTs that are connected to at least three IPs.

7 To filter the query results, see "Filtering Query Results in the Cardinality Tab" on page 77.

8 Click **OK** to save the changes you have made.

Filtering Query Results in the Cardinality Tab

Mercury Application Mapping enables you to define relationship conditions that allow the relationship's connecting node to be included in the TQL query results. In the **Cardinality** tab, you can create an expression defining that condition.

To filter your query results:

1 In the Editing pane, right-click the TQL node or relationship whose qualifier conditions you wish to define, and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Cardinality** tab.

3 In the **Cardinality** area, create an expression using the following:

Field	Description
Brackets { }	Click inside the **Brackets** box to display a list of brackets you can use to build more complex logical statements.
Criteria	Contains the attribute condition definition as defined in the Condition dialog box (for details, see step 3 in "Defining Attribute Conditions for Nodes and Relationships" on page 71).
And/Or	Click inside the **And/Or** field and select either **AND** or **OR** to link multiple conditions.

Note: If you choose not to create an expression, the default relationship AND is used between all relationships attached to the selected node.

4 To delete an attribute condition, select the attribute condition you want to delete and click the **Delete selected row** button.

5 To insert an attribute condition before a selected row, select the required row, and click the **Insert an attribute condition before a selected row** button to open the Condition dialog box.

6 To edit an attribute condition, select the attribute condition you want to edit and click the **Edit selected row** button to open the Condition dialog box. For details on how to define an attribute condition, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

7 To move a row up, select the row you want to move and click the **Move selected row up** button.

8 To move a row down, select the row you want to move and click the **Move selected row down** button.

9 Click **OK** to save the changes you have made.

Sample Relationship Conditions

The section contains examples of relationship conditions based on the sample TQL query shown below.

The following relationship conditions were defined for the TQL.

➤ [contained (Host, IP)] OR [member (Network, Host)] – means that the host must either have between two and four IPs OR be a member of the network.

➤ [contained (Host, IP)] AND [member (Network, Host)] – means that the host must have between two and four IPs AND ALSO be a member of the network.

Using the TQL Node Wizard

You can use the TQL Node Wizard to build a TQL query. The TQL Node Wizard allows you to:

➤ Add a node to a TQL query

➤ Add a relationship to a TQL query

➤ Define node conditions

➤ View all of the instances found for the TQL node in a table

This section includes:

➤ Adding a Node to the TQL Query – for details, see page 81

➤ Adding a Relationship to a TQL Query – for details, see page 83

➤ Setting TQL Node Definitions – for details, see page 85

➤ Showing TQL Node Instances – for details, see page 86

Adding a Node to the TQL Query

This section describes how to add a node to the TQL query.

To add a node to the TQL query:

1 In the required view, if the view is empty, drag a TQL node on to the Editing pane from the tree displayed in the Configuration Item Type pane.

2 Right-click the required TQL node and select **TQL Node Wizard** . The New TQL Node page opens, allowing you to to add a new TQL node to the query.

This page displays only the list of TQL nodes of the CIT type that have valid relationships to the selected (source) node.

To the right of each node, the number of CI instances that exist in the CMDB for that CIT type is displayed. The number of instances is only updated after you close the TQL Node Wizard and open it again.

The first node in the hierarchy tree is selected by default.

3 Select the **Show only CITs with instances** check box to display only the CITs of which there are instances in the CMDB.

Note: The **Show only CITs with instances** check box is selected by default.

---

4 (Optional) The **Element name** box contains the name of the selected node. By default, Mercury Application Mapping assigns the CIT type as the element's name. You can rename a TQL node in the **Element name** box giving it a unique label. This can be helpful when there is more than one node of the same CIT type in the TQL query.

5 The **<node> is required** check box defines the relationship cardinality (for details, see "Defining Relationship Cardinality" on page 76).

➤ Select the **<node> is required** check box for each node to include a minimum of one instance of the node at the other end of the relationship in the query results. Selecting the check box gives the relationship a cardinality value of 1*.

➤ Clear the **<node> is required** check box to give the relationship a cardinality value of 0*.

6 Click **Finish** to create the TQL query with the node you have added. To add a relationship, see "Adding a Relationship to a TQL Query" on page 83.

Adding a Relationship to a TQL Query

This section describes how to add a relationship to the TQL query.

To add a relationship to the TQL query:

1 In the New TQL Node page (for details, see "Adding a Node to the TQL Query" on page 81), select **Next** to open the New TQL Relationship page.

The Add Relationship dialog box contains the following fields:

Field	Description
Show only relationships with instances	Display only the relationships of which there are instances in the CMDB.
Node Label	The label of the source node. The default value is the CI Type display name. For details, see "Creating CITs" on page 384.
	Note: This field appears on the right hand side of the page.

Field	Description
Direction	The direction of the relationship that indicates which node is dependent on the other.
Node Label	The label of the node you selected in the New TQL Node page. **Note:** This field appears on the right hand side of the page.
Relationship Type	A valid relationship that defines the connection between the selected nodes. The field displays the relationship you selected in the node hierarchy tree.

2 Select the required direction of the relationship.

3 If required, select another valid relationship type to define the connection between the selected nodes.

Note: The relationship type that is selected by default is the highest in the hierarchy of the available, valid relationships.

---

4 Click **Finish** to create the TQL relationship you added to the query. The selected nodes are linked by the relationship you have selected. The direction of the relationship indicates which node is dependent on the other. To set node definitions, see "Setting TQL Node Definitions" on page 85.

Setting TQL Node Definitions

This section describes how to define TQL node definitions.

To define TQL node definitions:

 1 In the New TQL Relationship page (for details, see "Adding a Relationship to a TQL Query" on page 83), select **Next** to open the New TQL Node Condition page.
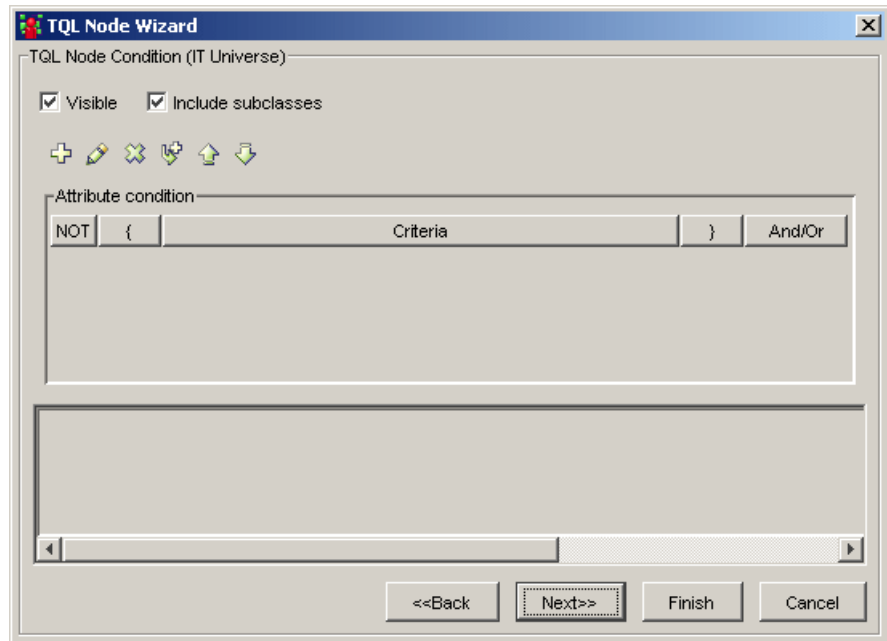
 2 To define attribute conditions for the node you added, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

 3 To filter query results in the **Attribute condition** area, see "Filtering Query Results in the Attribute Condition Tab" on page 73.

 4 Click **Finish** to save the node conditions you have defined. To display the TQL instances that were found, see "Showing TQL Node Instances" on page 86.

Showing TQL Node Instances

You can display all of the instances found for each TQL node in a table.

To display all of the instances found for each TQL node in a table:

 1 In the TQL Node Condition page (for details, see "Setting TQL Node Definitions" on page 85), select **Next** to open the Element Instances page. For details, see "Showing TQL Node Instances" on page 88.

 2 Click **Finish** to save your TQL query definitions.

Straightening a Relationship With an Angle

Mercury Application Mapping allows you to straighten a relationship that has angles.

To straighten a relationship with angles:

Right-click the relationship that has an angle and select **Straighten**.

Creating a Dependency Graph

You can create a graph that represents additional TQL query data related to a specific CI. The discovery pattern searches for the results from TQL query as well as the dependency graph definitions. The query recursively retrieves all the related CIs by a defined depth.

In the graph, you can define the relationship that is connected to a specific node. For example, if one of the nodes is a host, you can specify different relationships for NT, Router, and IP. You can also define attribute conditions for nodes. For details, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

The discovery pattern retrieves data that meets the criteria that is defined in the dependency graph.

This section contains the following topics:

➤ Defining a Dependency Graph – for details, see page 87

➤ Deleting a Dependency Definition – for details, see page 88

➤ Editing a Dependency Definition – for details, see page 88

Defining a Dependency Graph

This section describes how to define a dependency graph.

To define a dependency graph:

1 Right-click the required node and select **Add SubGraph** to open the Dependency List dialog box.

2 Click the **Add** button to open the Add Dependency dialog box.

3 From the **Node 1** list, select the desired node.

4 From the **Node 2** list, select the desired node.

5 To define attribute conditions for the node you selected, click the **Condition** button to the right of the node. For information on how to define attribute conditions, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

Note: The **Condition** button is enabled only after you selected a node from the node list.

---

6 From the **Relationship** list, select an available relationship connecting Node 1 to Node 2. The list of available relationships appears only after defining both nodes.

7 Click **OK** to save the dependency map definition. The dependency map definition appears in the Dependency List dialog box.

8 In the **Depth** box, select a number that determines how many steps (that is, how many connecting nodes) in the TQL pattern you want included in the discovery process. The default is 5.

9 Click **OK** to save the changes you have made.

Deleting a Dependency Definition

This section defines how to delete a dependency definition.

To delete a dependency definition:

 1 From the Dependency List dialog box, select the row you want to delete.

 2 Click the **Delete** button.

Editing a Dependency Definition

This section defines how to edit a dependency definition.

To edit a dependency definition:

 1 From the Dependency List dialog box, select the row you want to edit.

 2 Click the **Edit** button to open the Add Dependency dialog box.

 3 Edit the nodes and relationships as desired. For details, see "Defining a Dependency Graph" on page 87.

 4 To define attribute conditions for a node, click the **Condition** button to the right of the node.

 For details, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

 5 Click **OK** to save the changes you have made.

Showing TQL Node Instances

You can display all of the instances found for each TQL node in a table.

To display all of the instances found for each TQL node in a table:

 1 In the Editing pane, right-click the required TQL and select **Show element's instances** to open the Element instances dialog box.

 The list of nodes are divided into pages. The number at the bottom of the screen indicates which page is currently being displayed. For example, 2/4 means that it is the second out of four pages.

 2 To view other pages, use the left and right arrows.

3 To determine the number of node instances that appear on a page, do the following:

➤ Click the **Set bulk size** button to open the Set bulk size dialog box.

➤ Use the up and down arrows or type the number of node instances you want to appear on a page and click **OK**.

4 To update the table, click the **Refresh** button.

5 Click **OK** to save the settings you have defined.

Editing a TQL query

This section defines how to edit a TQL query.

To edit a TQL query:

1 In the View Explorer, right-click the TQL query you want to edit and select **Properties** to open the TQL Properties window.

2 Make the required changes. For details, see "Creating a TQL Query" on page 59.

3 Click **OK** to save the changes you have made.

Copying/Pasting TQL Nodes in the TQL Builder

You can copy and paste existing TQL nodes in the TQL Builder. The copied TQL nodes includes all the TQL node information. You can also copy and paste relationships provided that the TQL nodes to which the relationships are connected are also selected. A relationship on its own cannot be copied without its connecting TQL nodes.

To copy and paste an existing TQL node in the TQL Builder:

1 Select the TQL node you would like to copy. You can select multiple TQL nodes as well.

2 Right-click and select **Copy**.

3 Right-click again and select **Paste**. The TQL node(s) are added to the TQL.

Saving a Query

After you have defined the specific attribute conditions of the nodes and relationships in the query, the last step is to save the query to the CMDB. If you exit the TQL Builder without saving, a confirmation dialog box is displayed, requesting that you save your query.

To save the query:

Display the query to be saved in the editing pane, and click **Save** on the toolbar or select a different query or tab in the View Explorer. A message is displayed asking if you want to save the current query.

The query you created is saved to the CMDB.

Importing a TQL Query

You can import XML files that contain saved TQL queries to your TQL Builder. Use this option if you want to relocate TQL queries from one workstation to another.

To import a TQL query:

1 From the TQL Builder, open the **Map** menu and click **Import** to display the Import dialog box.

2 Locate the TQL query you want to import, select it and click **Import**. The imported TQL query is added to the query list in the View Explorer.

Exporting a Set of TQL Items

You can export a set of TQL items.

To export a set of TQL items:

1 Open the Manager whose TQL items you want to export. Then, open the **Map** menu and select **Export All** to open the Export All dialog box:
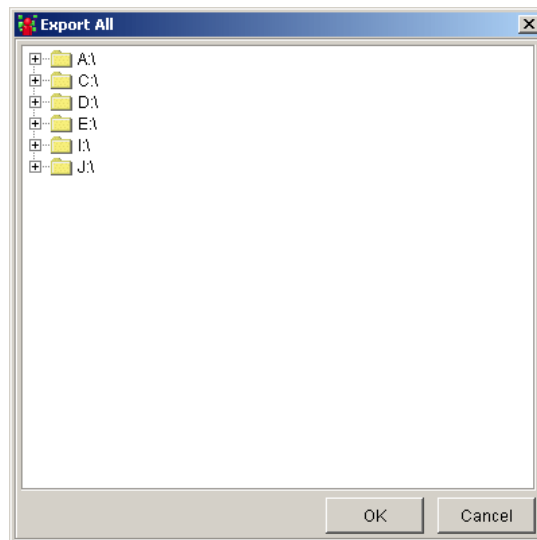
2 Browse and select the folder to which you want to export the TQL items displayed in the Manager.

3 Click **OK**.

5

Validation Restrictions

This chapter explains TQL validation restrictions.

About Validation Restrictions

For Correlation, Enrichment, and Report, and Discovery TQL types to be valid, they must comply with the following restrictions:

➤ **Internal Relationships** – A TQL must not contain internal relationships, that is, a relationship must not lead from a node to itself, as the following example illustrates:
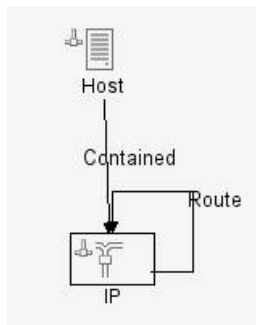
➤ **Circular Structure –** The TQL structure cannot be a closed circle, as shown in the following example:
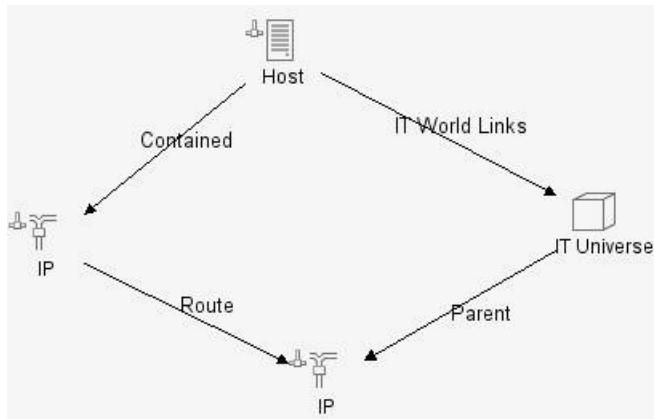
Note: There is one exception to the above rule: the TQL can be circular if one of its relationships is invisible. To define a node/relationship as invisible, see step 3 in "Defining TQL Nodes and Relationships" on page 69.

➤ **Separate Nodes and Groups –** All the TQL nodes must be linked to one another, meaning the TQL cannot contain separate nodes or groups, as the following example illustrates:
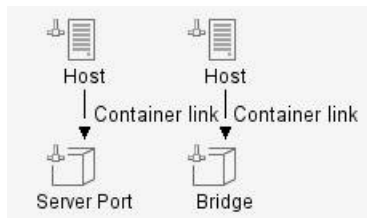
Correlation TQL Validation

In addition to the above restrictions, a Correlation TQL should also comply with the following restriction:

➤ **Number of Nodes –** The TQL must consist of at least two nodes.

Enrichment and Report TQL Validations

In addition to the above restrictions, Enrichment and Report TQLs should also comply with the following restriction:

Relationship definition – If a TQL node has a relationship, whose minimum limit is **0**, that is, one of the relationship ends does not necessarily have a node, the new logical node cannot be linked to it (since it may or may not exist in the TQL). For the procedure to define a relationship condition, see "Defining Attribute Conditions for Nodes and Relationships" on page 71.

Part III

Pattern Views and Folding Rules

6

Introduction to the Service View Manager

This chapter introduces the Service View Manager, which enables you to define service views for viewing the results of TQL queries.

This chapter describes:	On page:
About Creating Service Views	99
Working With Service View Manager	101
Service View Manager Shortcut Menus	102
Service View Manager Tooltips	106

About Creating Service Views

Mercury Application Mapping allows you to create a subset of your overall business environment, a service view, to define a view that displays only the information that is relevant to your organization's business needs.

The Service View Manager enables you to define pattern views, by which to populate your view. Pattern views are built from a TQL (Topology Query Language) query, which are a series of rules and definitions for displaying query results. By creating service views, you can define how the query results are displayed in the topology map, including nodes, groups and the number of map levels to be displayed.

The view only displays those CIs and relationships that satisfy the query definition. Each view is attached to a specific TQL query (if required, you can define multiple views for each query).

When the view is opened up or refreshed, the Service View Manager queries the CMDB for all elements that fit the query, and automatically updates the view with those elements. For more information on pattern views, see "Working with Pattern Views" on page 107. In addition, by creating view folders, you can determine how the views are arranged in the topology map.
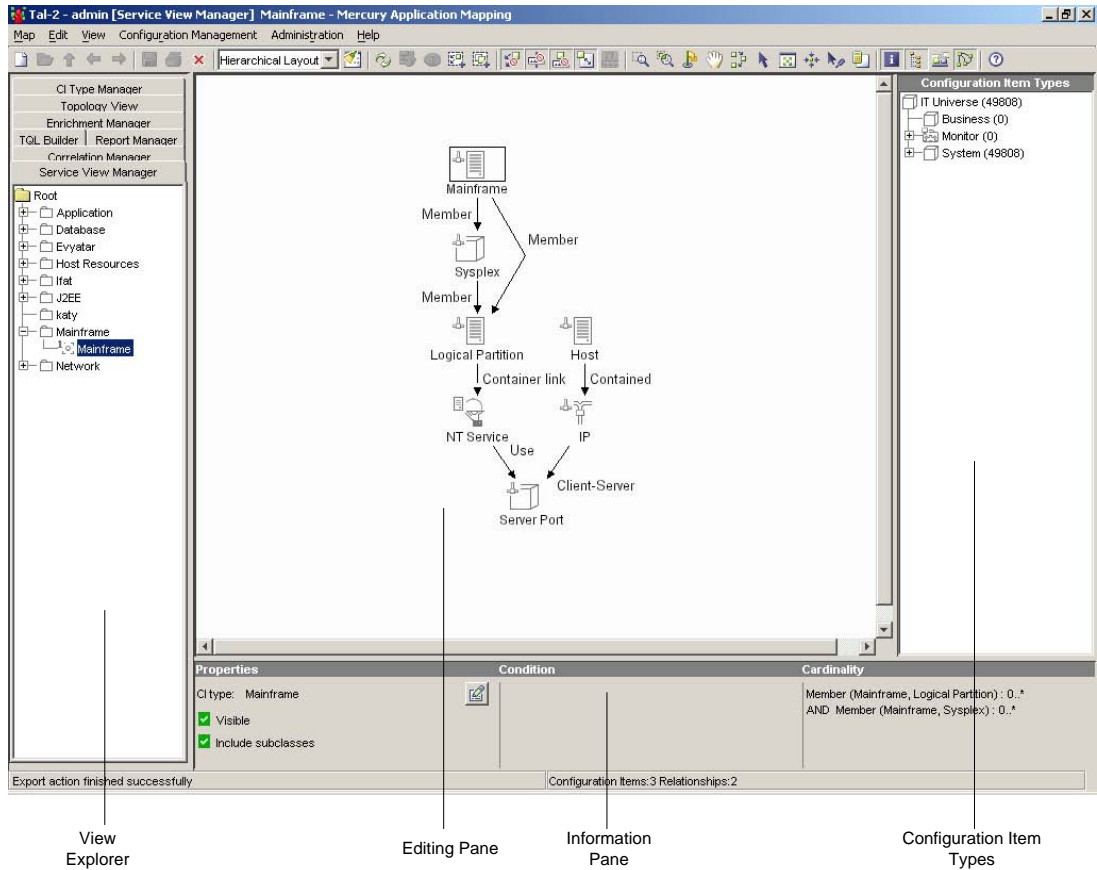
When the Service View Manager tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **View Explorer –** Displays a hierarchical tree structure of the views you have defined. Multiple views can be defined for each query, each of which can be used only with the query to which it is attached. The color of the view icon indicates the view status. For details, see "Creating a Pattern View" on page 110.

➤ **Editing Pane –** Displays the currently selected view, which consists of TQL nodes that are defined in the TQL query and the relationships between them.

➤ **Configuration Item Types –** Displays the the Configuration Item Types and the icons for each CIT, as defined by the administrator (for details, see "Assigning an Icon to a CIT" on page 389). By clicking and dragging CITs to the topology map and then defining the relationship between them, you can define a query and save it to the database. For example, you can drag the **NT** and **IP** CIs to the topology map and then define the relationship between them by adding relationships.

➤ **Information Pane –** Displays the attribute conditions you defined for the selected node and relationship. For details, see, "Setting TQL Node and Relationship Definitions" on page 124.

Working With Service View Manager

Use the Service View Manager to perform the following functions:

➤ Create new views, as well as duplicate and delete existing views.

➤ Determine how the nodes and relationships, which are defined in the query, are displayed in the topology map.

➤ Show or hide and delete selected nodes and relationships in the editing pane.

➤ Define the node's status factor (that is, its system-level significance which affects the status of the CI in the topology map.)

➤ Define a grouping option for the nodes displayed in the topology map.

➤ Define the attributes to display in the label of the selected node in the topology map.

➤ Add and delete organization rules that define the map levels to be displayed in the topology map.

➤ Save a view as an XML script. This can be used to relocate views from one workstation to another. (This can be performed providing the attached TQL query is also relocated.)

➤ Zoom in and out of the editing pane area to view the selected layer at different levels of magnification.

Service View Manager Shortcut Menus

The Service View Manager contains right-click shortcut menus, depending on your selection, as follows:

➤ "View Folder Shortcut Menu" on page 102

➤ "View Shortcut Menu" on page 103

➤ "Node Shortcut Menu" on page 104

➤ "Relationship Shortcut Menu" on page 105

View Folder Shortcut Menu

This is the same as the TQL Folder shortcut menu. For details, see "TQL Folder Shortcut Menu" on page 53.

View Shortcut Menu

The following table contains a brief description of each option in the menu displayed by right-clicking a selected view in the View Explorer:

Option	Description
New	Enables you to create a new view.
New Folder	Enables you to create a new view folder. These folders are displayed in the topology map and can assist you in organizing views according to criteria that are significant to your management needs, such as specific technologies, corporate procedures, internal workflows, company sites and so forth.
Add/Remove to Topology View	Enables you to add/remove views to the list in the topology map. When adding a view, a checkmark is displayed beside the selected view in the View Explorer. This option is especially useful for removing views from the topology map that are created by other users and may not be of interest.
Save As	Enables you to define a new name and description for a view. You can use this option to create a new view based on an existing one. For details, see "Creating a View by Copying an Existing View" on page 116.
Delete	Enables you to delete a selected view from the database.
Export	Enable you to save a view as an XML script. This option can be used to move views from one workstation to another, provided the attached TQL query is also relocated.
Properties	Displays the View Properties dialog box (which is the same as the Create New View dialog box). This option enables you to change the properties of the selected view (apart from the name of the view and the attached TQL). **Note:** You can also change the view's properties at a later stage from the topology map.
Relationship Rules	Displays the Relationship Rules dialog box, enabling you to define multiple relationship rules.

Node Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a node in the Topology pane:

Option	Description
TQL Node Wizard	Build a TQL query. For details, see "Using the TQL Node Wizard" on page 135.
Add Relationship	Define the connection between two nodes. For details, see "Adding Nodes and Relationships to a TQL Query" on page 116
TQL Node Definition	Define attribute conditions for nodes. For details, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.
Delete	Delete a node.
Add SubGraph	Create a graph that represents additional TQL query data related to a specific node. For details, see "Creating a Dependency Graph" on page 153.
Show element's instances	Display all of the instances found for each TQL node in a table. For details, see "Showing TQL Node Instances" on page 155.
View Node Definition	Display the View Node Definition dialog box, enabling you to determine how the nodes defined in the query are displayed in the topology map.
Notification	Display the Notification dialog box, enabling you to determine the condition under which notifications are sent. For details, see "Sending Notifications" on page 113.
Copy	Copy and paste a TQL node and its definitions into a view.
Add to Business Service	Define the creation of a business service item by view name and link it to different CIs. For details, see "Business Service Enrichment" on page 156.

Relationship Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a relationship in the Topology pane:

Option	Description
TQL Relationship Definition	Define attribute conditions for relationships. For details, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.
Delete	Delete a relationship.
Straighten	(Enabled only for relationships with angles on the Topology pane.) Straighten relationships with angles.
Add Rule	Define the organizational structure of nodes displayed in the topology map.
View Relationship Definition	Display the View Relationship Definition dialog box, enabling you to determine how the relationships defined in the query are displayed in the topology map.
Copy	Copy and paste a relationship and its definitions into the same view or into another view definition.
Edit Relationship Type	Edit an existing relationship. This option only appears if at least one of the following is true: ➤ You can change the direction the relationship. ➤ You can select a descendent of the existing relationship. ➤ The selected relationship is function relationship (Compound or Join)

Tip: Right-clicking an empty area in the topology map displays a shortcut menu with the same zoom and layout options available in the **View** menu. For details, see "Mercury Application Mapping Quick Tour" on page 17.

Service View Manager Tooltips

When the pointer is moved over a node in the topology map, a tooltip is displayed.

The tooltip contains the following information:

Option	Description
Element Name	The name of the node.
Visible	Whether the selected node is visible in the topology map.
Derived	Whether the nodes derived from the selected node are displayed in the topology map.
Tags	The tags that have been defined for the selected node and view.
Cardinality	Relevant for relationships. Cardinality defines how many nodes you expect to have at the other end of the relationship. For example, in a relationship between host and IP if the cardinality is 1:3, the TQL will retrieve only the hosts that are connected to 1 to 3 IPs.
Label Format	The format of the label displayed underneath the node.
Status Factor	The relative weight of the selected node, used for determining the status of the node.
Group node By	The type of group (CIT, TQL, node, or attribute) the node belongs to, if any.
Minimum in Group	The minimum number of matching nodes required to create the group to which the node belongs.
Group Label	The label of the group to which the node belongs.

7

Working with Pattern Views

This chapter describes how you create a pattern view. Pattern views enable you to query the CMDB by defining Topology Query Language (TQL) queries. The pattern view you create displays the results of the TQL queries.

This chapter describes:	On page:
Saving the Pattern View	159
Importing a Managed View	159

About Pattern Views

A pattern view enables you to define TQL queries that extract information from the CMDB that is relevant to your needs "Mercury Application Mapping Topology Query System" on page 9. You define a pattern view and then create or attach a TQL query, which is a series of rules for defining how the query results are displayed in the editing pane, including which nodes and relationships are to be displayed. Once a TQL query has been created, it persists in the system memory and continuously generates updated results.

About Defining Pattern Views

When you create a pattern view, you must first name the managed view and then define how you want the query results displayed, including organization (folding) rules that determine the number of map levels to be generated, attribute conditions for each node and relationship and how many CI instances you expect to have at the end of a relationship in your query result.

You can either define the pattern view by manually dragging and dropping the TQL nodes manually from the Configuration Item Types pane or use the TQL Node Wizard (for details, see "Using the TQL Node Wizard" on page 135).

Once you have finished defining how you want the TQL results displayed, you activate the TQL query to retrieve the requested data from the CMDB.

Pattern View Workflow

You create pattern views according to the following workflow:

➤ Define a new view.

For details, see "Creating a Pattern View" on page 110.

➤ Add nodes and relationships to the query.

For details, see "Adding Nodes and Relationships to a TQL Query" on page 116.

➤ Define node and relationship attribute conditions.

For details, see "Setting TQL Node and Relationship Definitions" on page 124.

➤ Add organization (folding) rules to relationships.

For details, see "Adding Organization (Folding) Rules to Relationships" on page 150.

➤ Define multiple relationship rules.

For details, see "Defining Multiple Relationship Rules" on page 151.

➤ Create a dependency graph.

For details, see "Creating a Dependency Graph" on page 153.

➤ Save the TQL query.

For details, see "Saving the Pattern View" on page 159.

Creating a Pattern View

The first step in creating a pattern view is to set the name and properties of the managed view in which you want to define the TQL query.

Note: TQL queries should only be created by administrators with an advanced knowledge of how the TQL system works.

This section includes:

Creating a View With the Service View Manager

You can create a view and attach it to a specify TQL query.

To create a pattern view:

1 Click the **Service View Manager** tab.

2 Click **New** on the toolbar or open the **Map** menu and click **New** to open the **View Properties** tab in the Create New View dialog box.

3 In the **View Name** box, enter a unique name for the new view.

4 (Optional) In the **Organization Name** box, enter the name of the organization related to the query.

5 (Optional) In the **Service Name** box, enter the name of the service related to the query.

6 (Optional) From the **ViewIcon** list, select a special icon for this view. The view is represented by this icon in the topology map.

7 In the **Integration Vendor** box, choose the integration application to which you want to connect Mercury Application Mapping.

8 (Optional) In the **Description** box, enter a description for the new view.

9 Select **Is Persistent** to define whether you want the selected view to always remain in the system memory. You should only use this option for views that are used frequently.

Note: Views that are not persistent are recalculated by the system only when they are selected for display. This may take some time, depending on the size of the view.

---

10 Select **Merge Identical Instances** to automatically display identical instances that appear on the same layer in this view merged into one symbol in the editing pane (for details, see "Merging Identical Instances" on page 112). This helps you remove redundant information and sharpen the view's focus.

11 In the **Send Notification on changes** area:

- Select **Added objects** if you want to be notified when CIs are added.

- Select **Removed objects** if you want to be notified when CIs are removed.

12 Click the **TQL Properties** tab to either create a new TQL or attach an existing one.

13 To create a new TQL, select **Create New TQL**.

14 In the **TQL Details** section, type a unique name for the TQL in the **Name** box.

15 From the **Priority** list, select a priority level for the new TQL query (Low, Medium, High, Express). This setting determines how often the query should be rerun automatically by the system to include updated information from the CMDB.

16 (Optional) In the **Description** box, enter a description of the TQL query.

17 If required, select **Is Persistent** to define whether you want the TQL to always remain in the system's memory. Use this option for TQLs that are used frequently rather than occasionally.

18 To attach a new TQL, select **Attach to existing TQL**.

19 In the **TQL Details** section, choose the TQL you want to attach from the **Name** list.

20 Repeat steps 15, 16, and 17.

21 Click the **View States** tab to define topology map states. For details, see "Defining Topology Map States" on page 115.

Merging Identical Instances

You can decide to have identical instances that appear on the same layer in this view automatically merged into one symbol on the editing pane. This helps you remove redundant information and sharpen the view's focus.

For example, suppose you want to discover and display hosts that are connected to other hosts in a server-client relationship, you can create the following TQL:
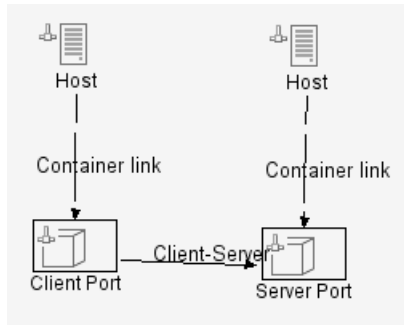
However, if both the server and the client are installed on the same host, one host appears twice.
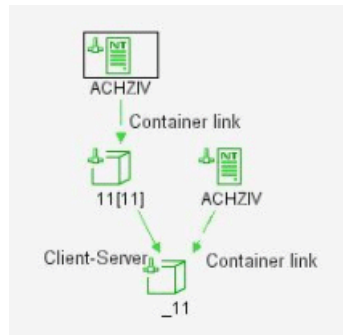
To overcome the redundant host display, you define that identical instances should be merged by selecting **Merge Identical Instances**:
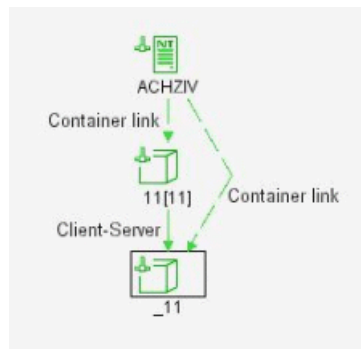
About Defining Topology Map States

To define the categories that are displayed on every map, you create a list of categories in the **Advanced** tab of the View Properties dialog box. By default, only the **change** state is predefined. To enable a specific map to receive and display additional categories, you add the categories to the **State Priority** list.

The **State Priority** list also determines the CI color according to the following rules:

➤ Mercury Application Mapping scans the CI state for each state that is on the list starting from the lowest number to the highest. The lowest number on the list has the highest priority.

For example, for the numbers 1, 2, 3, 4, the state that has the Priority value of 1 has the highest priority, the state that has the Priority value of 2 has the next highest priority, and so forth.

➤ If the CI state in the first state (the highest priority value) is not equal to **Normal**, the CI color is determined according to the state color that is defined in the Enumeration Manager dialog box. For details, see "About Enumerations and Lists" on page 67.

➤ If the CI state in the first state (the highest priority value) is equal to Normal, then Mercury Application Mapping checks the next highest state on the **Priorities** list (the state that has the next highest number) for a state that is not equal to Normal.

In the following example, the CI takes the color according to the Performance state because it is the first in the list of priorities whose severity is other than Normal. Security is a custom state.

Priority	State	Severity
1	Operation	0
2	Performance	5
3	Security	8

➤ In cases where there are two states of equal importance, you can assign them the same priority value providing they are defined based on the same Enumeration. In this case, the most critical value determines the CI color.

In the following example, the states Operation and Test have the same priority level and same Enumeration definition. The CI takes the color that represents the seventh severity in the Enumeration definition.

Priority	State	Severity
1	Operation	5
1	Test	7
2	Performance	8
3	Security	10

Defining Topology Map States

This section describes how to define states in the topology map

To define topology map states:

1 In the **View States** tab (for details, see step 21 in "Creating a View With the Service View Manager" on page 110).

2 Click the plus button to create a new row and add a new priority level.

Note: To remove a state priority, select it and click the remove button.

---

3 Click inside the **State** column and choose the state. The states that appear in the list are the states defined in the State Manager dialog box. For details, see "Defining Attribute States", in the *Mercury Application Mapping Administration Guide*.

4 Click inside the **Priority** column and enter a priority value.

Note: You can give the same priority value to two or more states only if they have the same enumeration definition. For details, see "About Enumerations and Lists" on page 67.

---

5 Repeat steps 2 to 4 to add another state to the list.

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

Creating a View by Copying an Existing View

You can create a new view by copying an existing view and modifying it.

To create a view by copying an existing view:

1 On the View Explorer, right-click the view you want to use as a basis for the new view. Click **Save As** to display the Save As dialog box.

2 Edit the view's definitions as required, as described in "Creating a Pattern View" on page 110.

3 Click **OK**.

View Manager Tooltips

Hold the pointer over a node or relationship to view its tooltip. The tooltips contain the definitions of the selected nodes and relationships, for example, the attribute conditions, as described in "Setting TQL Node and Relationship Definitions" on page 124.

Adding Nodes and Relationships to a TQL Query

Once you have created the pattern view, the next step is to add the nodes and relationships that define the query in the Create View page. The nodes represent the CITs, as defined in the CI Type Model. A relationship is a set of rules that defines the connection between two nodes. Relationships are defined one at a time for each pair of nodes in the query.

For a list of the relationships you can use to link two nodes and their definitions, see "Relationship Definitions" on page 543 .

This section contains the following topics:

➤ Adding Nodes and Relationships to Define a TQL Query – for details, see page 117

➤ Defining Function and Compound Relationships – for details, see page 119

➤ Sample Compound Relationship – for details, see page 121

➤ Deleting a Node or Relationship – for details, see page 123

➤ Editing a Relationship Type – for details, see page 123

Adding Nodes and Relationships to Define a TQL Query

This section describes how to add nodes and relationships to the pattern view to define the TQL query.

To add nodes and relationships to a define the TQL query:

1 From the tree displayed in the Configuration Item Type pane, click and drag one or more required TQL nodes on to the editing pane. These are the TQL nodes that are included in the query.

2 To add a relationship between two nodes:

➤ Select the required TQL node(s) by holding down CTRL and clicking the TQL nodes, right-click and select **Add Relationship**.

or

➤ Click the **Create Relationship** icon and draw a line between the required nodes.

The Add Relationship dialog box opens. It contains the following fields:

Field	Description
Node Label	The label of the selected nodes.
Relationship Direction	The direction of the relationship that indicates which node is dependent on the other.
Relationship Type	A valid relationship that defines the connection between the selected nodes.
Advanced	Define the connection between two nodes using either a function or compound relationship
	You can define the connection between nodes by adding either function or compound relationships. For details, see "Defining Function and Compound Relationships" on page 119.

 3 Select the required direction of the relationship.

 4 If necessary, click **Advanced** to select either function or compound relationships.

 5 Click **OK**. The selected nodes are linked by the relationship you have selected.

The direction of the relationship indicates which node is dependent on the other. The following example displays two hosts, an IP server, and an IP client that are linked to one another via a client/server connection.

Defining Function and Compound Relationships

You can define the connection between two nodes using either a function or compound relationship.

➤ A function relationship (**Virtual join link**) describes a logical connection between two nodes, which is not represented in the CMDB.

When you create a function relationship, you must define an attribute for each node, whose value will be used for comparison.

➤ A compound relationship (**Virtual compound link**) describes a connection that is composed of more than one relationship. If there are two nodes whose connection cannot be directly defined by one relationship, you can add several relationships, which act as stepping stones along the line, to create a connection.

To define function and compound relationships:

1 Select the nodes you want to connect with a function or compound relationship (for details on how to add nodes to a TQL query, see "Adding Nodes and Relationships to Define a TQL Query" on page 117).

2 Right-click and select **Add Relationship** to open the Add Relationship dialog box.

3 Click **Advanced**.

4 From the **Relationship** list, select **Function Relationship**.

5 To add a **Virtual join link**, do the following:

➤ Select **Virtual - Join link**.

➤ In the **Join Definition** area, click the **Add an attribute definition** button to open the Conditions dialog box.

➤ From the **Attribute** lists, choose the two required attributes.

➤ From the **Operator** list, select one of the following operators:

• Select **Equal** if you want the system to check whether the two selected attributes are equal.

• Select **Not equal** if you want the system to check whether the two selected attributes are not equal.

119

➤ Click **OK** to save your definition. The definition appears in the **Join Definition** area.

➤ To delete a definition, select the required definition and click the **Delete selected row** button.

➤ To edit a definition, select the required definition and click the **Edits selected row** button to open the Conditions dialog box. Edit the definition as described in step 5.

➤ Click **OK** to save your changes.

6 To add a **Virtual compound link**, do the following:

➤ Select **Virtual - Compound link.**

➤ In the **Compound Definition** area, click the **Add** button to open the Add dependency dialog box.

➤ From the **Source** list, select the source node.

➤ From the **Target** list, select the target node.

➤ From the **Relationship** list, select an available relationship connecting the two nodes. The list of available relationships appears only after defining both nodes.

➤ From the **Relationship Direction** list, select the required direction.

Note: You can create as many virtual compound definitions as necessary.

➤ Click **OK** to save the compound definition. The compound definition appears in the Compound definition dialog box.

Note: Each node represents a step in the chain leading from one node to the other. Ensure that <**node 1**> of the first row and <**node 2**> of the last row are the two selected nodes in the TQL query.

➤ To delete a definition, select the required definition and click the **Delete** button.

➤ To edit a definition, select the required definition and click the **Edit** button to open the Add dependency dialog box. Edit the definition as described in step 6.

➤ In the **Depth** box, select a number that determines how many steps (that is, how many connecting nodes) in the TQL pattern you want included in the discovery process. The default is 5.

7 Click **OK** to save your definitions.

Sample Compound Relationship

This section contains a sample compound relationship comprising several nodes and their connections, displayed in two different views:

➤ A normal view in which each node in the compound relationship was dragged and dropped from the Configuration Item Types pane. (For details, see"Adding Nodes and Relationships to Define a TQL Query" on page 117.)

The normal view result is displayed below:

➤ A compound view in which the compound relationship was defined by using the **Compound Definition** area of the Add Relationship dialog box, as illustrated below. (For details, see "Defining Function and Compound Relationships" on page 119.)

The following compound view is created:

The compound view result is as follows:

Deleting a Node or Relationship

This section describes how to delete a node or relationship.

To delete a node or relationship:

Right-click the TQL node or relationship you want to delete and select **Delete**.

Editing a Relationship Type

This section describes how to edit an existing relationship. This option only appears if at least one of the following is true:

➤ You can change the direction the relationship.

➤ You can select a descendent of the existing relationship.

➤ The selected relationship is function relationship (Compound or Join)

To edit a relationship type:

1 Right-click the required relationship and select **Edit Relationship Type**. The Edit Relationship Type dialog box opens.

2 Make the required changes. For details, see "Adding Nodes and Relationships to Define a TQL Query" on page 117 and "Defining Function and Compound Relationships" on page 119.

3 Click **OK** to save your changes.

Setting TQL Node and Relationship Definitions

After you have added the TQL nodes and relationships required for your query, you can define their specific attribute conditions.

This section includes:

➤ Defining TQL Nodes and Relationships – for details, see page 124

➤ Changing the CI Type of the TQL Node – for details, see page 126

➤ Defining Attribute Conditions for Nodes and Relationships – for details, see page 126

➤ Filtering Query Results in the Attribute Condition Tab – for details, see page 129

➤ Defining Qualifier Conditions – for details, see page 130

➤ Defining Relationship Cardinality – for details, see page 131

➤ Filtering Query Results in the Cardinality Tab – for details, see page 132

Defining TQL Nodes and Relationships

This section describes how to define TQL nodes and relationships.

To define TQL nodes and relationships:

1 In the editing pane, right-click the TQL node or relationship whose attribute conditions you wish to define, and select **TQL Node Definition** or **TQL Relationship Definition** to open the TQL Node/Relationship Definition dialog box.

Note: Alternatively, you can select the required node/relationship, and click the **Go to Edit Dialog** button in the Information pane to open the TQL Node/Relationship Definition dialog box.

2 (Optional) The **Element name** box contains the name of the selected node/relationship. By default, Mercury Application Mapping assigns the CIT type as the element's name. You can rename a TQL node in the **Element name** box giving it a unique label. This can be helpful when there is more than one node of the same CIT type in the TQL query.

3 To display the selected TQL node in the editing pane, select **Visible**. When **Visible** is cleared, an invisible box appears to the right of the selected TQL node in the editing pane:

Any query results pertaining to that TQL node are not displayed in the topology view. This can be useful when certain relationships or TQL nodes are required to build the query but are not needed in the results. For example, NTs are connected to the network by defining specific IPs, but in any query results you may want to view the IP elements only and not the NT elements

4 To display both the selected CI and its descendents in the topology map, select **Include subclasses**.

Changing the CI Type of the TQL Node

You can change the CI type of the TQL node after you have created a TQL.

To change the CI type of a TQL node after creating a TQL query:

 1 In the editing pane, right-click the TQL node whose CI type you want to change. and select **TQL Node Definition** to open the TQL Node Definition dialog box.

 2 Click the **Change CI type** button to open the Change CI type dialog box.

 3 Select the required CI type and click **OK**.

Note: You can only change the CI type to a type from one of the CIT's subclasses, if one exists.

Defining Attribute Conditions for Nodes and Relationships

This section describes how to define attribute conditions for nodes and relationships.

To define attribute conditions:

 1 In the editing pane, right-click the TQL node or relationship whose attribute conditions you wish to define, and select **TQL Node Definition** or **TQL Relationship Definition** to open the TQL Node/Relationship Definition dialog box.

 2 In the TQL Node/Relationship Definition dialog box, click the **Attribute Conditions** tab.

 3 Click the **Add an attribute condition** button to open the Condition dialog box.

 4 From the **Attribute Name** list, choose the required attribute.

 5 From the **Operation** list, select the required operation.

The operations in the **Operation** list are:

Operation	Description
Equal	Checks whether the attribute value is equal to the value specified in the **Value Comparison** box.
Not equal	Checks whether the attribute value is not equal to the value specified in the **Value Comparison** box.
Is null	Checks whether the attribute value is null.
Like	Uses a wildcard (%). Use **Like** when you are not sure of the complete name of what you are looking for.
In	Displays only the instances where this attribute value equals one of the selected values. For example, for CIs that have an operational state that equals Warning(2) and Warning(3), select the operation **In**, and select both Warning(2) and Warning(3) from the **Value comparison** list.
Equal ignore case	Checks whether the attribute value is equal to the value specified in the value box regardless of the case.
Greater	Checks whether the attribute value is greater than the value specified in the **Value comparison** box.
Great than or equal	Checks whether the attribute value is greater than or equal to the value specified in the **Value comparison** box.
Less	Checks whether the attribute value is less than the value specified in the **Value comparison** box.
Less than or equal	Checks whether the attribute value is less than or equal to the value specified in the **Value comparison** box.
Changed during	(Displayed only when you select the **Create Time** attribute.) Displays only the instances that changed during the period specified in the **Value comparison** box.

Operation	Description
Unchanged during	(Displayed only when you select the **Create Time** attribute.) Displays only the instances that did not change during the period specified in the **Value comparison** box.
In list	Equals one of the elements that appears in the list. For example: 320,4445,3483.

Note: When using the **Not equal** or **Not like** operations, query results do not include data from the CI instances that were not assigned a value. For example, let us assume that your system contains three hosts. Host1 was assigned the value A, Host2 was assigned the value B, and Host3 was not assigned a value. If you created a query in which you want to retrieve all hosts that are **Not equal** to A, the results of your query only include Host2 because Host3 was not assigned a value.

6 If you want to make a value comparison, do the following:

➤ Select **Value**.

➤ In the **Value Comparison** box, enter or select the value of the attribute.

Note: If you use a string in the **Value** box, it must appear in single quotes, as follows: 'YOKO', 'JASON'.

7 If you want to create a condition in which you compare two attributes, do the following:

➤ Select **Attribute**.

➤ In the **Attribute Comparison** box, select the required attribute from the **Attribute Comparison** list.

8 To filter the query results, see "Filtering Query Results in the Attribute Condition Tab" below.

9 Click **OK** to save the changes you have made and close the Condition dialog box.

Filtering Query Results in the Attribute Condition Tab

In the **Attribute Condition** tab, you can create an expression defining conditions that restrict the number of nodes that appear in the query.
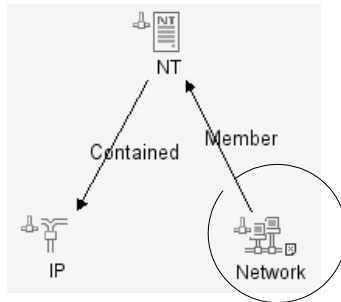
To filter your query results:

1 In the editing pane, right-click the TQL node or relationship whose attribute conditions you wish to define, and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Attribute Condition** tab.

3 In the **Condition** area, create an expression using the following:

Field	Description
NOT	Select **NOT** if you want the condition statement to do the opposite of what is defined.
	Note: If you select **NOT**, query results do not include data from the CI instances that were not assigned a value. For example, let us assume that your system contains three hosts. Host1 was assigned the value A, Host2 was assigned the value B and Host3 was not assigned a value. If you created a query in which you want to retrieve all hosts that are equal to A and selected **NOT**, the results of your query only include Host2 because Host3 was not assigned a value.
Brackets `{ }`	Click inside the **Brackets** box to display a list of brackets you can use to build more complex logical statements.
Criteria	Contains the attribute condition definition as defined in the Condition dialog box (for details, see step 3 in "Defining Attribute Conditions for Nodes and Relationships" on page 126).
And/Or	Click inside the **And/Or** field and select either **And** or **Or** to link multiple conditions.

4 To delete an attribute condition, select the attribute condition you want to delete and click the **Delete selected row** button.

5 To insert an attribute condition before a selected row, select the required row, and click the **Insert an attribute condition before a selected row** button to open the Condition dialog box.

6 To edit an attribute condition, select the attribute condition you want to edit and click the **Edit selected row** button to open the Condition dialog box. For details on how to define an attribute condition, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.

7 To move a row up, select the row you want to move and click the **Move selected row up** button.

8 To move a row down, select the row you want to move and click the **Move selected row down** button.

9 Click **OK** to save the changes you have made.

Defining Qualifier Conditions

This section describes how to define qualifier conditions for the selected node/relationship. For example, you can use a qualifier to define a CIT as abstract, meaning you cannot create instances from it.

To define qualifier conditions:

1 In the editing pane, right-click the TQL node or relationship whose qualifier conditions you wish to define, and select **TQL Node Definition** or **TQL Relationship Definition** to open the TQL Node/Relationship Definition dialog box.

2 Click the **Qualifier condition** tab.

3 Select the required qualifier(s) according to the following:

- ➤ **NOTIFY_CHANGES** – For internal use only

- ➤ **BLE_LINK_CLASS** – For internal use only

- ➤ **ABSTRACT_CLASS** – You cannot create instances of this CIT

- ➤ **ITU_HIDDEN_CLASS** – For internal use only

- ➤ **CONTAINER** – Relevant for relationships only. Represents containment between two CIs.

➤ **HIDDEN_CLASS** – Does not appear anywhere in the application.

➤ **READ_ONLY_CLASS** – For internal use only

Note: You can hold the CTRL key down to make multiple selections.

➤ Click **OK** to save your definitions.

Defining Relationship Cardinality

In this section, you define how many CI instances you expect to have at the end of a relationship in your query result.

To define relationship cardinality:

1 In the editing pane, right-click the TQL node whose conditions you wish to define, and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Cardinality** tab.

Note: This tab only appears when a node is selected.

---

3 Click the **Add an attribute condition** button to open the Relationship Cardinality dialog box.

4 In the **Node** box, select the required relationship to which the selected node is attached. The list will contain all the relationships that are linked to the selected node.

In the following steps, you define the lower and upper limits for including the node at the other end of the relationship in the query results.

5 In the **Min** box, enter the value that defines the lower limits for including the node at the other end of the relationship in the query results.

6 In the **Max** box, enter the value that defines the upper limits for including the node at the other end of the relationship in the query results. You can use an asterisk (*****) in the **Max** box to represent an infinite value.

For example, if <end_1> is IP and <end_2> is NT, entering 1 in the **Min** box and asterisk (*****) in the **Max** box instructs the system to retrieve only those IPs that are connected to at least one NT. (The asterisk indicates an infinite value.) Entering 3 in the **Min** box and asterisk (*****) in the **Max** box instructs the system to retrieve only those NTs that are connected to at least three IPs.

7 To filter the query results, see "Filtering Query Results in the Cardinality Tab" on page 132.

8 Click **OK** to save the changes you have made.

Filtering Query Results in the Cardinality Tab

Mercury Application Mapping enables you to define relationship conditions that allow the relationship's connecting node to be included in the TQL query results. In the **Cardinality** tab, you can create an expression defining that condition.

To filter your query results:

1 In the editing pane, right-click the TQL node or relationship whose qualifier conditions you wish to define, and select **TQL Node Definition** to open the TQL Node Definition dialog box.

2 Click the **Cardinality** tab.

3 In the **Cardinality** area, create an expression using the following:

Field	Description
Brackets `{` `}`	Click inside the **Brackets** box to display a list of brackets you can use to build more complex logical statements.
Criteria	Contains the attribute condition definition as defined in the Condition dialog box (for details, see step 3 in "Defining Attribute Conditions for Nodes and Relationships" on page 126).
And/Or	Click inside the **And/Or** field and select either **And** or **Or** to link multiple conditions.

4 To delete an attribute condition, select the attribute condition you want to delete and click the **Delete selected row** button.

5 To insert an attribute condition before a selected row, select the required row, and click the **Insert an attribute condition before a selected row** button to open the Condition dialog box.

6 To edit an attribute condition, select the attribute condition you want to edit and click the **Edit selected row** button to open the Condition dialog box. For details on how to define an attribute condition, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.

7 To move a row up, select the row you want to move and click the **Move selected row up** button.

8 To move a row down, select the row you want to move and click the **Move selected row down** button.

9 Click **OK** to save the changes you have made.

Sample Relationship Conditions

This section contains an example of relationship conditions based on the sample TQL query shown below.
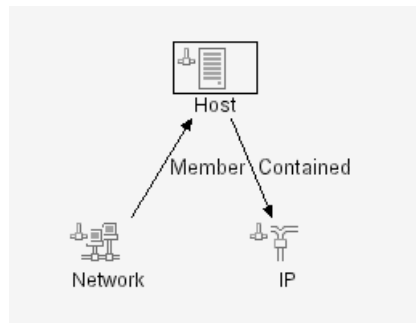
The following relationship conditions were defined for the TQL:

In the Relationship Cardinality dialog box, the following was defined (see "Defining Relationship Cardinality" on page 131):

➤ Member – **Min**: 1, **Max**: *

➤ Contained – **Min**: 2, **Max**: 4

The definitions appear in the **Cardinality** section, as displayed below.
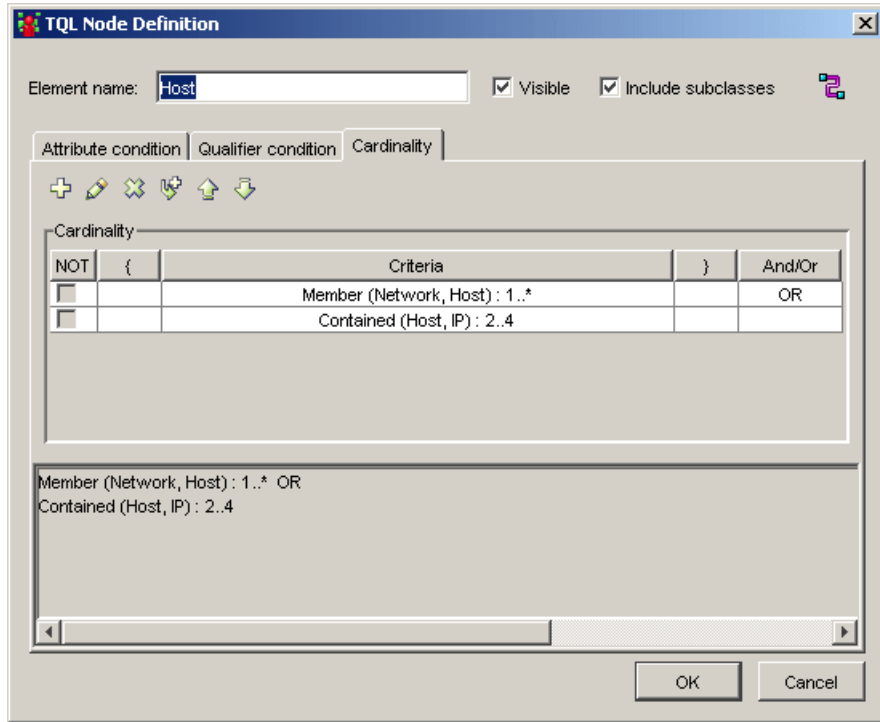
➤ [contained (Host, IP)] OR [member (Network, Host)] – means that the host must either have between two and four IPs OR be a member of the network.

➤ [contained (Host, IP)] AND [member (Network, Host)] – means that the host must have between two and four IPs AND ALSO be a member of the network.

Using the TQL Node Wizard

You can use the TQL Node Wizard to build a TQL query. The TQL Node Wizard allows you to:

➤ Add a node to a TQL query

➤ Add a relationship to a TQL query

➤ Define node conditions

➤ View all of the instances found for the TQL node in a table

This section includes:

➤ Adding a Node to the TQL Query – for details, see page 135

➤ Adding a Relationship to a TQL Query – for details, see page 138

➤ Setting TQL Node Definitions – for details, see page 140

➤ Showing TQL Node Instances – for details, see page 141

Adding a Node to the TQL Query

This section describes how to add a node to the TQL query.

To add a node to the TQL query:

1 In the required view, if the view is empty, drag a TQL node on to the editing pane from the tree displayed in the Configuration Item Type pane.

2 Right-click the required TQL node and select **TQL Node Wizard**. The New TQL Node page opens, allowing you to add a new TQL node to the query.

This page displays only the list of TQL nodes of the CIT type that have valid relationships to the selected (source) node.

To the right of each node, the number of CI instances that exist in the CMDB for that CIT type is displayed. The number of instances is only updated after you close the TQL Node Wizard and open it again.

The first node in the hierarchy tree is selected by default.

3 Select the **Show only CITs with instances** check box to display only the CITs of which there are instances in the CMDB.

Note: The **Show only CITs with instances** check box is selected by default.

4 (Optional) The **Element name** box contains the name of the selected node. By default, Mercury Application Mapping assigns the CIT type as the element's name. You can rename a TQL node in the **Element name** box giving it a unique label. This can be helpful when there is more than one node of the same CIT type in the TQL query.

5 The **<node> is required** check box defines the relationship cardinality (for details, see "Defining Relationship Cardinality" on page 131).

➤ Select the **<node> is required** check box for each node to include a minimum of one instance of the node at the other end of the relationship in the query results. Selecting the check box gives the relationship a cardinality value of 1*.

➤ Clear the **<node> is required** check box to give the relationship a cardinality value of 0*.

6 Click **Finish** to create the TQL query with the node you have added. To add a relationship, see "Adding a Relationship to a TQL Query" on page 138.

Adding a Relationship to a TQL Query

This section describes how to add a relationship to the TQL query.

To add a relationship to the TQL query:

 1 In the New TQL Node page (for details, see "Adding a Node to the TQL Query" on page 135), select **Next** to open the New TQL Relationship page.

The Add Relationship dialog box contains the following fields:

Field	Description
Show only relationships with instances	Display only the relationships of which there are instances in the CMDB.
Node Label	The label of the source node. The default value is the CI Type display name. For details, see "Creating CITs" on page 384. **Note:** This field appears on the right hand side of the page.
Direction	The direction of the relationship that indicates which node is dependent on the other.
Node Label	The label of the node you selected in the New TQL Node page. **Note:** This field appears on the right hand side of the page.
Relationship Type	A valid relationship that defines the connection between the selected nodes. The field displays the relationship you selected in the node hierarchy tree.

 2 Select the required direction of the relationship.

 3 If required, select another valid relationship type to define the connection between the selected nodes.

Note: The relationship type that is selected by default is the highest in the hierarchy of the available, valid relationships.

---

 4 Click **Finish** to create the TQL relationship you added to the query. The selected nodes are linked by the relationship you have selected. The direction of the relationship indicates which node is dependent on the other. To set node definitions, see "Setting TQL Node Definitions" on page 140.

Setting TQL Node Definitions

This section describes how to define TQL node definitions.

To define TQL node definitions:

1 In the New TQL Relationship page (for details, see "Adding a Relationship to a TQL Query" on page 138), select **Next** to open the New TQL Node Condition page.
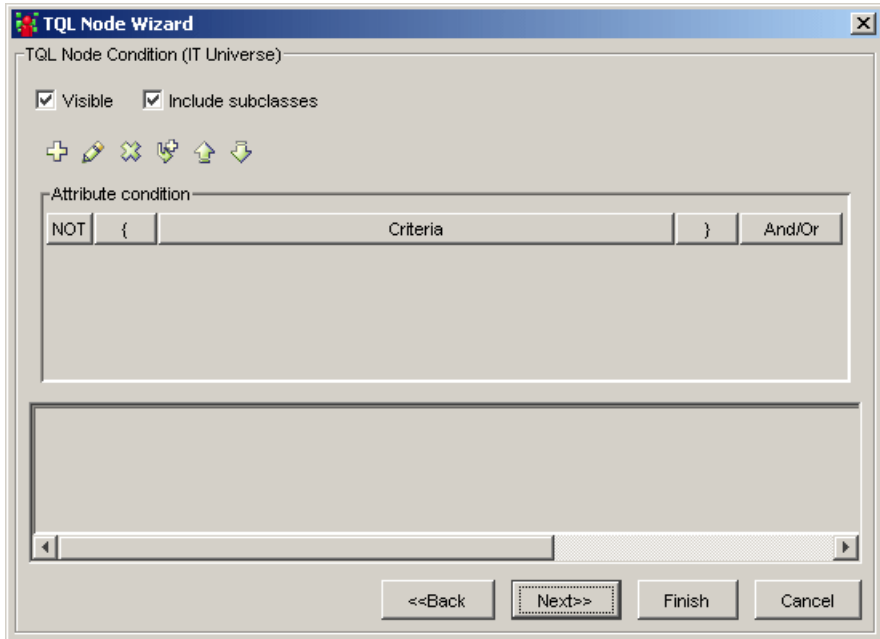
2 To define attribute conditions for the node you added, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.

3 To filter query results in the **Attribute condition** area, see "Filtering Query Results in the Attribute Condition Tab" on page 129.

4 Click **Finish** to save the node conditions you have defined. To display the TQL instances that were found, see "Showing TQL Node Instances" on page 141.

Showing TQL Node Instances

You can display all of the instances found for each TQL node in a table.

To display all of the instances found for each TQL node in a table:

1 In the TQL Node Condition page (for details, see "Setting TQL Node Definitions" on page 140), select **Next** to open the Element Instances page. For details, see "Showing TQL Node Instances" on page 155.

2 Click **Finish** to save your TQL query definitions.

Specifying the CI Status Factor

The status factor of a CI is a value (in percent) that defines the significance of the CI status. This is useful for reducing the significance of CIs whose status can fluctuate on a regular basis, such as processes, which continuously change. By reducing the significance of such CIs in the system, events are prevented from generating an overflow of critical messages. For example, if the state of an event is 5 and the status weight of the process is defined at 40%, then the CI status is 2 (5*40/100).

To specify the status factor of a CI:

1 In the editing pane, right-click the node for which you want to define the factor status and select **View Node Definition**. The View Node Definition dialog box opens.

2 Click the **Status Factor** tab.

3 In the **Status Factor** area:

> ➤ In the **Status Weight** box, enter a value (in percentage) to define the significance of the CI's status. This is useful for reducing the significance of CIs whose status can fluctuate on a regular basis, such as processes and CPU. By reducing the significance of such CIs in the system, you can prevent events from generating an overflow of critical messages.

> For example, if the state of an event is **5** and the status weight of the process is defined at **40%**, then the CI status is **2 (5*40/100)**. For details on status, see "Introduction to the Topology View" on page 163

> ➤ In the **Compound Status Weight** box, enter a value (in percentages) to define the significance of the CI's compound status (for details, see compound status on page 550).

4 In the **Propagation** area:

> ➤ (Optional) Select **Propagate Node Status** to determine whether the status of the node is inherited by its upper view's layers.

> ➤ (Optional) Select **Propagate Node Blink** to determine whether the blinking attribute of the node (indication for its unacknowledged state) is inherited by its upper view's layers.

> ➤ (Optional) If you selected the **Propagate Node Status** option above, select whether the propagation should be the maximum (**Max**) or average status of the CI (**Average**) in the **Propagate Function** list.

5 Click **OK** to save the settings you have defined.

Grouping CIs

Mercury Application Management enables you to group the CIs in a pattern view according to defined criteria, for example, by CIT or by a particular attribute. You can apply the grouping option you defined to the layer of your view to organize the way the view is presented throughout Mercury Application Management. For more details, see "Grouping Options (Pattern Views Only)" on page 25.

You can define whether different grouping definitions should be used for the topology and tree presentations of the view. To define different grouping options for topology and tree presentations, you select the **Use different folding/grouping definitions** option in the view properties dialog box (for details, see "Using the Different Folding/Grouping Definitions Option" on page 22).

You can apply grouping options for the overall view, or different grouping options can be defined to apply for the topology presentations and tree presentations of the view.

This section has the following topics:

➤ "Grouping CIs in the Topology Map" on page 143

➤ "Grouping By Attribute Properties" on page 147

➤ "Nested Groups" on page 148

Grouping CIs in the Topology Map

This section describes how to group CIs in the topology map.

To group CIs in the topology map:

1 In the editing pane, right-click the node for which you want to create groups according to selected criteria and select **View Node Definition**. The View Node Definition dialog box appears.

2 Click the **Group By** tab to group the node's instances according to selected criteria.

3 To define a group of nodes for this view, click **Add** to display the Group Configuration Items dialog box.

4 Define a group of nodes for this view by selecting one of the following options:

➤ **Configuration Item Type** – Group CIs in the topology map by their CIT.

➤ **TQL Node** – Groups all CIs of the same CIT that fulfill the node definition criteria. By default, these CIs appear in the topology map under a folder, containing the name of the selected TQL node and an automatically added identification number.

➤ **Attribute** – Groups CIs from multiple CITs in the topology map, according to the attribute selected from the list. A CI can be a member of multiple groups. The **Attribute** option enables you to group CIs from different CITs. The **Configuration Type Item** and **TQL Node** options are used to group CIs from the same CIT. For details, see "Grouping By Attribute Properties" on page 147).

5 (Optional) In the **Minimum in group** box, enter the minimum number of matching CIs required to create the group. By default, there must be at least one CI in the group.

Note: If you use the Minimum in group option, click the **Rebuild** button in the toolbar to view updated results.

 6 (Optional) In the **Group label** box, enter the label of the group.

 7 You can create nested groups (using the **operation state** attribute), that is, a CI group inside another CI group. This option provides you with additional fine-tuning of the TQL results that are displayed in the topology map (for more details, see "Nested Groups" on page 148).

 8 To nest a group, create one group and return to the **Group By** dialog box (for more details, see "Nested Groups" on page 148). Click the **Add** button again and create another group. When you finish creating the second group, the **Group By** tab should look similar to the following figure:
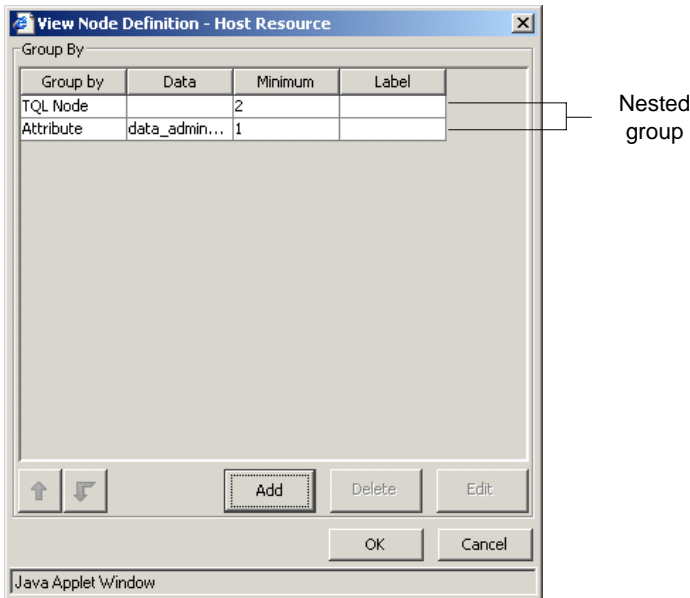
You can repeat this procedure to create as many groups as needed.
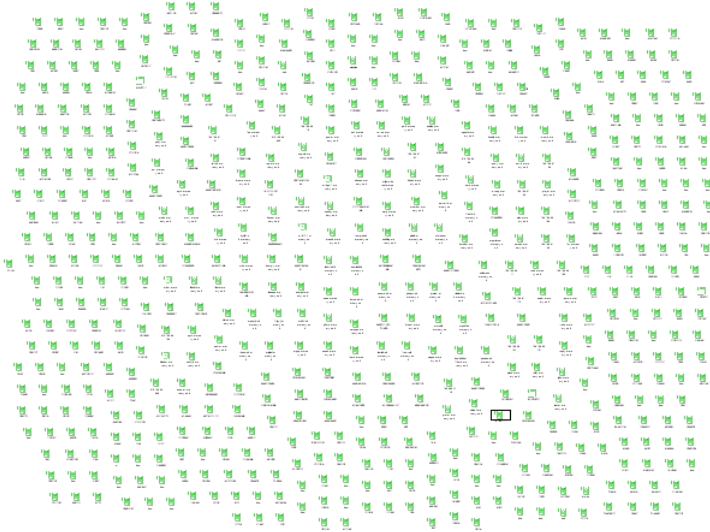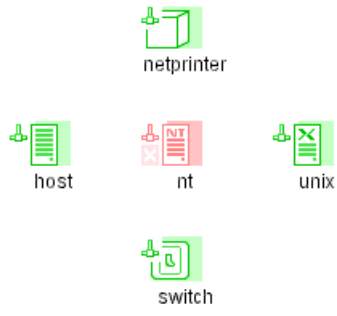
 9 To change their nesting order by using the arrow buttons at the bottom of the View Node Definition dialog box.

10 To delete a group, select the required group and click **Delete**.

11 To edit a group:

➤ Select the required group and click **Edit**.

➤ Repeat steps 3 to 5 in "Grouping CIs in the Topology Map" on page 143.

12 Click **OK** to save the node's definitions.

13 Click **OK** to save the definitions of the selected node.

This following contains examples showing how the host grouping affects the display in the topology map:

	Non-grouped hosts
	Grouped hosts

Grouping By Attribute Properties

You can group nodes from multiple CITs in the topology map, according to the following attributes:

➤ **Mask to group by:**

Enter a regular expression, as follows:

➤ In the first field, enter the regular expression pattern. This is the structure of the selected attribute.

➤ In the second field, enter the group number. This is the part of the regular expression pattern to focus on when creating the group.

For example, enter a regular expression to define the IP address (aa.yy.zz.mm), as follows:

To:	In the First Field Enter:	In the Second Field Enter:
Group by zz	(.*[.].*[.])(.*)([.].*)	2
Group by yy	(.*[.])(.*)([.].*[.].*)	2
Group by aa	(.*)([.].*[.].*[.].*)	1
Group by mm	(.*[.].*[.].*[.])(.*)	2

You can also group by the first or last letter of any attribute:

For example, enter a regular expression to group the nodes by the first or last letter(s) of the selected attribute as follows:

To:	In the First Field Enter:	In the Second Field Enter:
Group by the first letter	(.)(.*)	1
Group by the last letter	(.*)(.)	2
Group by the first two letters	(..)(.*)	1
Group by the last two letters	(.*)(..)	2

➤ **Group List** – Define the specific groups of the selected attribute to be created. For example, to display two groups of IP addresses that end between 0 and 100, enter 0-50, 51-100. Any IPs that fall outside this group are placed automatically into the **Others** group.

Nested Groups

You can create nested groups (a node group inside another node group) to help fine-tune the TQL results that are displayed in the topology map.

For example, after grouping host nodes by their CIT, you can use their **operation state** attribute to create more focused groups.

In the topology map, nested grouping is displayed as follows:

➤ The first layer shows host nodes grouped by their CIT:

netprinter

host   nt   unix

switch
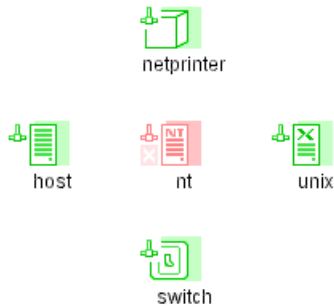
➤ Drill down to the next layer to display host nodes grouped by **Oper State**:

Critical   Normal

➤ Drill down to the next layer to display host nodes belonging to a certain host CIT and with the same **Oper State**:
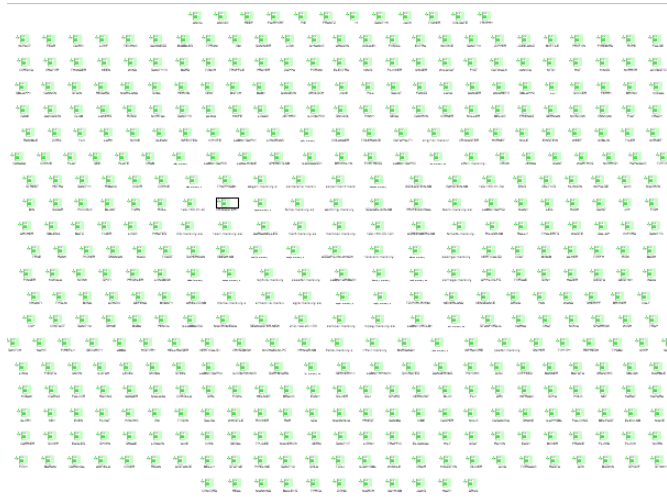
Straightening a Relationship With an Angle

Mercury Application Mapping allows you to straighten a relationship that has angles.

To straighten a relationship with angles:

Right-click the relationship that has an angle and select **Straighten**.

Adding Organization (Folding) Rules to Relationships

After setting the node and relationship definitions of the view, you can add organization (folding) rules to selected relationships. This enables you to define the organizational structure of CIs in the topology map by displaying selected CIs at different levels. When no folding rules are defined, the topology view displays all CIs included in the query results on one level by default. For example, if the query results include hosts and networks, both CIs are displayed on the same level in the topology map.

If an icon appears with a shadow in the topology map, it means that there is an additional CI layer beneath it, created by a parent or child folding rule.

You can define whether different folding definitions should be used for the To define different folding options for topology and tree presentations, you select the **Use different folding/grouping definitions** option in the view properties dialog box (for details, see "Using the Different Folding/Grouping Definitions Option" on page 22).

To add folding rules:

1 In the editing pane, right-click the relationship for which you want to define a folding rule.

2 Select **Add Rule**.

3 Select one of the following options to define the organizational structure between the two CI types connected by the selected relationship:

➤ **Parent** – Select this option to display Node1 as the parent of Node2. That is, double-clicking the Node1 element on one map level displays a new map level with Node2 elements.

➤ **Right Sibling** – Select this option to display Node1 elements wherever Node2 elements are displayed. That is, both CIs are displayed on the same map level.

➤ **Left Sibling** – Select this option to display Node2 elements wherever Node1 elements are displayed. That is, both CIs are displayed on the same map level.

➤ **Child –** Select this option to display Node1 as the child of Node2. That is, double-clicking a Node2 element on one map level displays a new map level with Node1 elements.

➤ **None –** Do not define any folding rules for this relationship.

Defining Multiple Relationship Rules

As an alternative to adding folding rules (for details, see "Adding Organization (Folding) Rules to Relationships" on page 150), you can choose to define multiple relationship rules.

Whereas you define organizational rules to define the structure of two nodes connected to a specific relationship, you define relationship rules to define the structure between CITs in a view.

For example, if one of the view's nodes is a host, you can specify different relationships for NT, UNIX, router, and so forth. By assigning different folding rules to different CITs, views can encompass and represent additional query data.

Note: If both folding and relationship rules have been defined, the definitions for the folding rules override the definitions for relationship rules.

This section includes the following topics:

➤ Defining Multiple Relationship Rules – for details, see page 152

➤ Editing Relationship Rules – for details, see page 153

➤ Deleting Relationship Rules – for details, see page 153

Defining Multiple Relationship Rules

Note: You can define sibling and child rules for the same relationship between two CITs. However, you cannot define parent and child rules for the same relationship, or left sibling and right sibling rules.

To define multiple relationship rules:

1 In the View Explorer, right-click the required view for which you want to define relationship rules and select **Relationship Rules** to open the Relationship Rules dialog box.
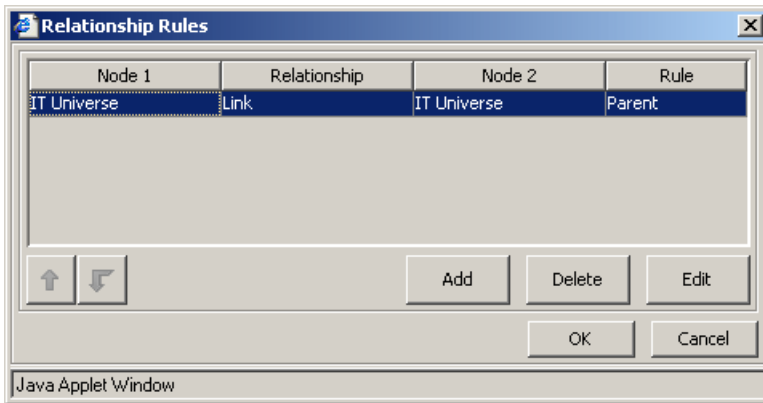
2 Click **Add** to display the Add Rule dialog box.

3 Select the required Node1 node in the **Node1** list.

4 Select the required Node2 node in the **Node2** list.

5 From the **Relationship** list, select an available relationship connecting the Node1 node to the Node2 node. (Enabled only after you defined both CITs.)

6 From the **Rule** list, select a folding rule: **Parent**, **Right Sibling**, **Left Sibling**, or **Child**. For details, see "Adding Organization (Folding) Rules to Relationships" on page 150.

7 Click **OK** to save the settings you have defined. The new rule is added to the Relationship Rules dialog box.

8 Repeat steps 2 to 7 to add another rule.

9 Click **OK**.

Editing Relationship Rules

This section describes how to edit an existing relationship rule.

To edit a relationship rule:

1 In the Relationship Rules dialog box, select the rule you want to edit and click **Edit** to display the Add Rule dialog box.

2 Make the required changes.

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

Deleting Relationship Rules

This section describes how to delete an existing relationship rule.

To delete a relationship rule:

In the Relationship Rules dialog box, select the rule you want to delete and click **Delete**.

Creating a Dependency Graph

You can create a graph that represents additional TQL query data related to a specific CI. The discovery pattern searches for the results from TQL query as well as the dependency graph definitions. The query recursively retrieves all the related CIs by a defined depth.

In the graph, you can define the relationship that is connected to a specific node. For example, if one of the nodes is a host, you can specify different relationships for NT, Router, and IP. You can also define attribute conditions for nodes. For details, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.

The discovery pattern retrieves data that meets the criteria that is defined in the dependency graph.

This section contains the following topics:

➤ Defining a Dependency Graph – for details, see page 154

➤ Deleting a Dependency Definition – for details, see page 155

➤ Editing a Dependency Definition – for details, see page 155

Defining a Dependency Graph

This section describes how to define a dependency graph.

To define a dependency graph:

1 Right-click the required node and select **Add SubGraph** to open the Dependency List dialog box.

2 Click the **Add** button to open the Add Dependency dialog box.

3 From the **Node 1** list, select the desired node.

4 From the **Node 2** list, select the desired node.

5 To define attribute conditions for the node you selected, click the **Condition** button to the right of the node. For information on how to define attribute conditions, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.

Note: The **Condition** button is enabled only after you selected a node from the node list.

---

6 From the **Relationship** list, select an available relationship connecting Node 1 to Node 2. The list of available relationships appears only after defining both nodes.

7 Click **OK** to save the dependency map definition. The dependency map definition appears in the Dependency List dialog box.

8 In the **Depth** box, select a number that determines how many steps (that is, how many connecting nodes) in the TQL pattern you want included in the discovery process. The default is 5.

9 Click **OK** to save the changes you have made.

Deleting a Dependency Definition

This section defines how to delete a dependency definition.

To delete a dependency definition:

1 From the Dependency List dialog box, select the row you want to delete.

2 Click the **Delete** button.

Editing a Dependency Definition

This section defines how to edit a dependency definition.

To edit a dependency definition:

1 From the Dependency List dialog box, select the row you want to edit.

2 Click the **Edit** button to open the Add Dependency dialog box.

3 Edit the nodes and relationships as desired. For details, see "Defining a Dependency Graph" on page 154.

4 To define attribute conditions for a node, click the **Condition** button to the right of the node.

For details, see "Defining Attribute Conditions for Nodes and Relationships" on page 126.

5 Click **OK** to save the changes you have made.

Showing TQL Node Instances

You can display all of the instances found for each TQL node in a table.

To display all of the instances found for each TQL node in a table:

1 In the editing pane, right-click the required TQL and select **Show element's instances** to open the Element instances dialog box.

The list of nodes are divided into pages. The number at the bottom of the screen indicates which page is currently being displayed. For example, 2/4 means that it is the second out of four pages.

2 To view other pages, use the left and right arrows.

3 To determine the number of node instances that appear on a page, do the following:

➤ Click the **Set bulk size** button to open the Set bulk size dialog box.

➤ Use the up and down arrows or type the number of node instances you want to appear on a page and click **OK**.

4 To update the table, click the **Refresh** button.

5 Click **OK** to save the settings you have defined.

Copying and Pasting TQL Nodes and Relationships

You can copy and paste TQL nodes and relationships and their definitions into the same view or another view definition.

To copy and paste TQL nodes and relationships:

1 Select the node(s)/relationship(s) you want to copy into the same view or another view. To select more that one item, hold down CTRL and click the required TQL nodes/relationships.

2 Right-click the selected items and click **Copy**.

3 In the view where you want to paste the selected item(s), right-click and select **Paste**.

Business Service Enrichment

Mercury Application Mapping allows you to define the creation of a business service item by view name and link it to different CIs. You can use this feature to show which hosts are serving a specific business service or create correlation rules that reflect problems at the business level.

You first create the required view, define the business view item in that view and then create the TQL query to see the results.

The TQL results in the topology map of the Topology View manager display the business service item and all the CIs linked to it. The name of the business service item gets its name of the view from which you defined the business service item. The business service item and the linked CIs are connected by **Application Links**.

By default, you can only apply this option to hosts. You can change the default by adding other CITs to the **appilogConfig.properties** file, located in \<**Mercury Application Mapping root directory**>**root****lib****server**\.

This section contains the following topics:

➤ Creating a Business Service View – for details, see page 157

Removing a Business Service Item from a Business Service View – for details, see page 159

Creating a Business Service View

This section describes how to create a business service view.

To create a business service view:

1 Create the required view and define the Business Service item. For example, in the following view called **Oracle**, right-click **Host** and select **Add to Business View**.
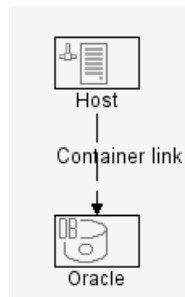
2 To see the results of the view, create the required TQL. For example, you can create a TQL using the **BusinessService** node and the **Application Link** as shown in the following example. For details about how to create TQL queries, see "Using the TQL Node Wizard" on page 135.
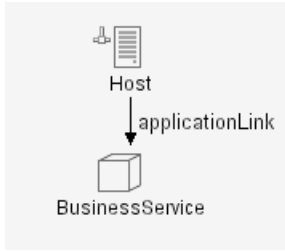
3 In this example, if you defined Host in your view as a business service item, a new business service item named **Oracle** is created and all the Hosts queried in this TQL are linked to an **Oracle** business service item by an Application Link.

The following example shows the business service view, which contains the **Oracle** business service item and the different CIs that are linked to it.
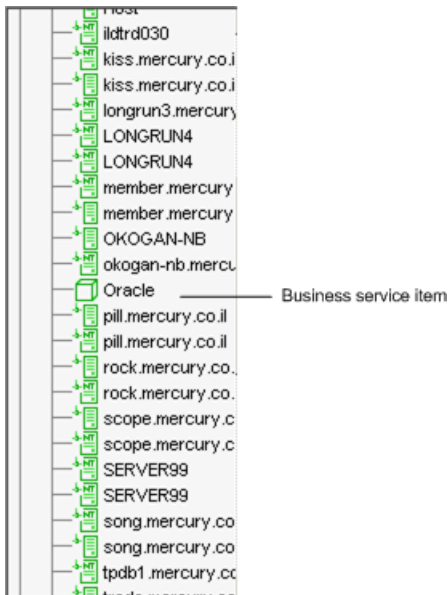
Removing a Business Service Item from a Business Service View

This section describes how to remove a business service item from a business service view.

To remove a business service item from a business service view:

Right-click the CI you defined as the business service item and select **Remove from Business Service**. The business service item is removed from the business service view.

Saving the Pattern View

Once you have defined your pattern view, the last step is to save the TQL query to the CMDB so it can retrieve the data based on the definitions you have set. If you exit the Service View Manager without saving, a confirmation message is displayed, requesting that you save your changes.

To save the pattern view:

Click the **Save** button on the toolbar.

Importing a Managed View

You can import XML files that contain saved views to the Service View Manager. Use this option if you want to relocate managed views from one workstation to another.

Note: Before you import a managed view, you must verify that its attached TQL query appears in the **Attached TQL** list in the New dialog box. If the attached query does not exist in your query list, importing fails.

To import a managed view:

1 From the Service View Manager, select the **Map** menu and click **Import** to open the Import dialog box.

2 Locate the managed view you want to import, select it and click **Import**. The imported managed view is added to the view list in the View Explorer.

Part IV

Viewing Queries with Topology View

8

Introduction to the Topology View

This chapter introduces the Topology View, which enables you to view multi-level maps displaying the results of the query based on a selected view. In addition, you can manage the events occurring in the system, at the CI level and at higher, top-view levels, enabling you to locate problems occurring in different areas of your IT infrastructure.

This chapter describes:	On page:
About the Topology View	164
Working with Topology View	165
Understanding Topology View Concepts	166
Topology View Shortcut Menus	167

About the Topology View

The topology view is displayed by selecting the Topology View tab in the View Explorer, and enables you to view the results of your TQL queries based on the view defined with the View Manager. The maps displayed in the topology map may be displayed in multiple layers, depending on whether organization rules are defined in the view, as described in Part III, "Pattern Views and Folding Rules." The topology map also enables you to view events associated with the CIs displayed in each map.

When the Topology View tab is selected, the Mercury Application Mapping main window is divided, as follows:

➤ **View Explorer –** Displays a hierarchical tree structure of the map CIs based on a selected view. You can navigate between map layers by either selecting layers from the View Explorer or by drilling down. For example, by selecting the view of a network in the View Explorer, you can drill down to view the CIs belonging to the selected network at the host level. Selecting a CI in the View Explorer simultaneously selects it in both the Map and Information panes. The color of the selected view's icon is determined by its severity level.

➤ **Filtered By pane –** Contains CI groups that can be displayed or hidden in the topology map.

➤ **Topology Map pane –** Displays the results of the view currently selected in the View Explorer, and consists of CIs that are defined in the TQL query and the relationships between them. You can change the way the maps are displayed by selecting one of the view layouts (Hierarchical, Symmetric, Manual, Orthogonal). For details, see "Improving Views" on page 197.

➤ **Information pane –** Displays events that are related to the CIs in the selected view. This pane includes the following tabs: **Configuration Item Browser**, **Configuration Item Events**, and **View Events**. For details, see "The Event Tabs in the Information Pane" on page 238.

Working with Topology View

Use the topology map to perform the following functions:

➤ Display the views created with the View Manager.

➤ Show or hide selected CIs, relationships or views.

➤ Edit labels of selected CIs in the Topology Map pane.

➤ Display the neighbors of a selected CI, which are connected to the CI using the various relationship types.

➤ Display and edit the CI's attributes, according to your access rights.

➤ Acknowledge events relating to the selected CIs and to the CIs displayed in the levels below the selected CIs (the compound CIs).

➤ Show or hide tags, which represent logical groupings of the CIs and relationships that are defined in the query.

➤ Display the selected CI's events and compound events.

➤ Zoom in and out of the Topology Map pane to view the selected layer at different levels of magnification.

➤ Display the view in one of a number of different layouts. If required, you can create your own manual layout view and save it to the database.

➤ Add a background map for navigational purposes, when using the manual layout display.

Understanding Topology View Concepts

The CIs and relationships that are represented in the Topology Map pane have unique topology data in addition to their basic attributes. This topology data is based on, and produced according to, the way these CIs and relationships are located and defined in Mercury Application Mapping views. You can view this data through the tooltip and browser, as described in the following sections.

This section includes:

➤ "Topology View Tooltip" on page 166
➤ "Configuration Item Browser" on page 167

Topology View Tooltip

When you point to a CI in the Topology Map pane, a tooltip is displayed that contains the CI statuses in the selected categories of the view:

Label:	**CPU0**		
Configuration Item Type:	**cpu**		

Category	Status	Correlation	Is new
operation	✅ Normal	✅ Normal	

The tooltip contains the following information:

➤ **Label** – The label of the selected CI.

➤ **Configuration Item Type** – The CIT of the selected CI as defined in the Configuration Item Type Model.

➤ **Category** – The category name.

➤ **Status** – The severity level of the category.

➤ **Correlation** – Indicates the severity of the CI that is affected by a correlation rule (in the corresponding category).

➤ **Is new** – Indicates that the CI has a new event that is neither acknowledged nor suppressed.

Configuration Item Browser

The **Configuration Item Browser** tab lists all the CIs of the selected view that are currently displayed in the Topology Map pane. (To display the **Configuration Item Browser** tab, select **View** > **Information Browsers** > **Configuration Item Browser.**) In addition to the information displayed in the tooltip, the **Configuration Item Browser** tab contains the following information:

➤ **Symbol** – The icon defined for the CIT of the selected CI.

➤ **No. of Children** – The number of children and relationships that exist at all map levels below the selected level. Double-click the selected value to view the map level below the current level.

Note: Selecting a CI in the **Configuration Item Browser** also selects it in both the Topology View pane and the View Explorer pane.

Topology View Shortcut Menus

The topology view contains different shortcut menus, depending on your selection in the Topology Map pane.

This section includes:

➤ "View Shortcut Menu" on page 168

➤ "CI Shortcut Menu" on page 170

➤ "Multiple CIs Shortcut Menu" on page 173

➤ "Relationship Shortcut Menu" on page 174

View Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a selected view in the View Explorer, Topology Map or Information panes:

Note: If you open the shortcut menu from the Information pane, you have one additional option – creating and displaying charts. For details, see Chapter 13, "Generating Charts."

Option	Description
Hide	Enables you to hide the selected view.
Properties	Displays the View Properties dialog box. This option enables you to see and edit the properties of the selected view (apart from the view name and the attached TQL).
Rebuild	Refreshes the CIs and relationships displayed in the selected view.
Snapshot	Enables you to view and to save snapshots.
Reports	Select one of the following options: **Asset Report**–lists all the CIs in a selected view and their attribute values **Dependency Report** –lists all the relationships of the type dependency that connect between servers in the selected view or between the servers in the CMDB and the servers within the selected view **Event Report** –displays the events for the CIs in the selected view

Option	Description
Events	(These options are enabled for parent CIs only.) Select one of the following options: ➤ **Compound Event Browser –** Displays the Event Browser dialog box that enables you to view active events related to CIs in lower maps that are connected to the selected view. Through this dialog box you can perform other actions. ➤ **Compound Event Log –** Displays the Event Log dialog box that enables you to view all the raw events related to CIs in lower maps that are connected to the selected view. Through this dialog box you can also create a chart.
Deactivate/Activate View	Toggles between showing and hiding the view's CIs. The CIs are not deleted.

CI Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a selected CI in the Topology Map, View Explorer, or Information panes.

Option	Description
Hide	Enables you to hide the selected CI. By hiding one or more CIs, you can focus on the CIs that are of interest to you. This option is useful for display purposes and can assist you in creating reports and diagrams, which are based on the current display of the topology map.
Show CI Attributes F9	Displays the CI Attributes dialog box, enabling you to view and edit CI attributes, according to your access rights.
Note	Enables you to create personal notes that include comments, and attach them to CIs or events in the topology map.
Delete	Enables you to delete the selected CI from the view and from the CMDB. **Note:** When you delete a parent CI, the selected CI and its children are removed from the view, but only the selected CI is removed from the database.
Get Related CIs	Enables you to display the interdependencies of a selected CI (that is, its neighbors) in different contexts: a specific layer, a specific view or the entire CMDB. For details on the Get Related CIs options, see "Displaying Interdependent CIs" on page 205.
Insert CI	Displays the Insert Configuration Item dialog box, enabling you to manually add new CIs to the database and to the view. For details, see "Building the Topology Database", in *Mercury Application Mapping Administration Guide*.
Reports	Enables you to choose reports for events, changes, and related CIs. Displays the appropriate dialog box, enabling you to create and display reports. For details, see Chapter 12, "Generating Map Reports."

Option	Description
Correlation	Select one of the following options: ➤ **Show Impact–**Displays the Impact Correlation Rules dialog box, enabling you to select a correlation rule, and then display in a new window all the CIs that are affected by the selected CI. If only one correlation rule is defined for this CI, the Impact window is directly displayed. ➤ **Problem Isolation–**Displays the Root-cause Correlation Rules dialog box, enabling you to select a correlation rule, and then display in a new window the root-cause CI(s) that affected the selected CI. If only one correlation rule is defined for this CI, the Correlation window is directly displayed.
Events	Select one of the following options: ➤ **Event Browser–**Displays the Event Browser dialog box, enabling you to view the events related to the selected CI. ➤ **Event Log–**Displays the raw event history of the selected CI.

Option	Description
Actions (*cont'd*)	Select one of the following options:
	➤ **Set Admin State** – Enables you to manually change the Admin State of a CI, so that it is no longer a managed CI in the system. For example, by setting the Admin State of a router to **Testing**, the router has no influence on other CIs in the system, and the other CIs have no influence on it. The default Admin State is **Managed**. Other options include: **Disabled, Restricted, Unknown,** and **Unmanaged**.
	➤ **Send Event** – Displays the Send Event dialog box, to enable you to define the event category that affects the CI, according to your access rights. For details on defining the event category, see "Sending an Event" on page 252.
	➤ **Clear All Events** – Enables you to clear the CI's state and return it to **Normal**. For details on clearing the CI's state, see "Clearing All Events" on page 252.
	➤ **Add Discovery Pattern** – Displays the Add Discovery Pattern dialog box, enabling you to manually invoke a Discovery Pattern for the selected CI. You can use this option to discover additional information about the CI through one of the available Discovery Patterns. The Add Discovery Pattern dialog box also allows you to edit the parameters of the Discovery Patterns and to customize them to your needs.
	➤ **Remove Discovery Pattern** – Displays the Remove Discovery Pattern dialog box, enabling you to manually remove a Discovery Pattern from the selected CI.
Ping	(This option is enabled for host and IP CIs only.) Enables you to ping the selected CI.

Multiple CIs Shortcut Menu

The table below contains a brief description of each option in the shortcut menu, displayed by selecting several CIs and right-clicking them in the Topology Map pane (for details, see "Selecting Multiple CIs" on page 43).

Note: Use Windows conventions—Shift+arrow key, Ctrl+arrow key—to select adjacent (for the Explorer tree only) or nonadjacent CIs.

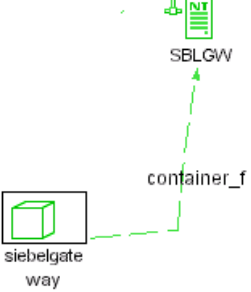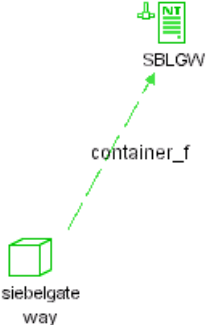You cannot select multiple CI from different view layers.

Option	Description
Hide	Enables you to hide the selected CI. By hiding several CIs, you can focus on the CI(s) that are of interest to you. This option is useful for display purposes and can assist you in creating reports and diagrams, which are based on the current display of the Topology Map pane.
Delete	Enables you to delete the selected CIs.
Actions	Select one of the following options: ➤ **Add CI to Discovery Pattern** – Displays the Add Discovery Pattern dialog box, enabling you to manually invoke a Discovery Pattern for the selected CI. You can use this option to discover additional information about the CI through one of the available Discovery Patterns. The Add Discovery Pattern dialog box also allows you to edit the parameters of the Discovery Patterns and to customize them to your needs. ➤ **Remove CI from Discovery Pattern** – Displays the Remove CI from Discovery Pattern dialog box.
Insert Relationship	(Enabled for two selected CIs only)Display the Insert Relationship dialog box, enabling administrators to manually add relationships between CIs that are saved to the database. Adding a relationship to the database triggers the Discovery Methods to begin collecting the required information regarding the new relationship.

Relationship Shortcut Menu

This section contains a brief description of each option in the shortcut menu displayed by right-clicking a relationship in the topology map pane:

Option	Description
Hide	Enables you to hide the selected relationship.
Show CI Attributes	Displays the Attributes dialog box, enabling you to view and edit relationship attributes, according to your access rights. For details, see "Viewing and Editing CI Attributes" on page 201.
Delete	Enables you to delete the selected relationship from the view and from the CMDB.
Correlation	Select one of the following options: ➤ **Show Impact** – Displays the Impact Correlation Rules dialog box, enabling you to select a correlation rule, and then display in a new window all the CIs that are affected by the selected relationship. If only one correlation rule is defined for this relationship, the Impact window is directly displayed. For further details, see "Showing a Correlation Impact" on page 212. ➤ **Problem Isolation** – Displays the Root-cause Correlation Rules dialog box, enabling you to select a correlation rule, and then display in a new window the root-cause CI(s) that affected the selected relationship. If only one correlation rule is defined for this CI, the Correlation window is directly displayed. For further details, see "Showing Root Cause CIs" on page 214.
Events	**Event Browser** – Displays the Event Browser dialog box, enabling you to view the events related to the selected relationship. For details, see "Event Browser" on page 248. ➤

Option	Description
Actions	Select one of the following options: ➤ **Set Admin State** – Enables you to manually change the Admin State of a relationship, so that it is no longer a managed CI in the system. For example, by setting the Admin State of a router to Testing, the router has no influence on other CIs in the system, and the other CIs have no influence on it. The default Admin State is Managed. Other options include: Disabled, Restricted, Unknown, and Unmanaged. ➤ **Send Event** – Displays the Set State dialog box, enabling you to determine the relationship's state, according to your access rights. ➤ **Clear Events** – Enables you to clear the relationship's Operational State and return it to Normal.

Option	Description
Straighten	(Enabled only for relationships with angles on the Topology Map pane.) Enables you to straighten relationships with angles, such as the one that appears in the following example: To straighten the relationship, right-click the blue square in the middle of the angle and select **Straight**. The relationship is straightened, as seen below:

9

Using the Topology Map

This chapter explains how to use the Topology View manager.

This chapter describes:	On page:
About Using the Topology Map	178
Understanding Layer Layouts	179
Deactivating and Activating Views	196
Improving Views	197
Viewing and Editing CI Attributes	201
Viewing Layers with a Large Number of CIs	203
Displaying Interdependent CIs	205
Showing a Correlation Impact	212
Showing Root Cause CIs	214
Using Meta Relationships	216
Saving a Topology Map	218
Printing the Topology Map	220
Sorting, Hiding, and Displaying Columns	222
Searching for CIs	223

About Using the Topology Map

The Topology Map pane enables you to navigate between different map levels (from the top level down to the CI level) to focus on specific areas of the network, according to your requirements. In addition, the Topology Map pane enables you to view the attributes of any displayed CI and add a background graphic to the map.

Navigating in Maps

When a view is selected in the View Explorer, the map for that view is displayed in the topology map pane. The icon displayed for each CI represented in the map is determined by its CIT, as defined in the CI Type Model, and includes the label defined for that CIT in the Service View Manager.

If the CI has a shadow, you can double-click the CI to drill down to a lower map level, created by a child organizational rule set up in the Service View Manager.

Host_Resources

For example, if the IP CI is defined as a child of the host CI in the Service View Manager, you can double-click any host icon in the topology map pane to display a map of IPs connected with that host.

To drill down from a selected parent CI, you can also use the Open option on the **Map** menu, or the **Open** button on the toolbar.

Right-click a CI and select from the shortcut menu one of the **Get Related CIs** options, to display a map containing the interdependencies of the selected CI in different contexts: a specific layer, a specific view or the entire Mercury Universal CMDB. Selecting the **Get Related CIs** > **Layer** option is useful when working with maps that display a great number of CIs, as it enables you to focus on a particular portion of the map while removing unneeded details. The other options enable you to locate and manage the interdependencies that exit between the selected CI and other CIs in your system, as described in "Displaying Interdependent CIs" on page 205.

Understanding Layer Layouts

The **Layer Setup** option enables you to accurately define the layout of view's layers in a way that helps you achieve the clearest presentation possible. By assigning logical positions for the nodes and relationships that appear on the layer, you can fine-tune your view's visualization, and thus facilitate the understanding and monitoring of your managed data.

This option applies only to nodes and relationships that are displayed in the topology map. The definitions that are set in the dialog boxes apply to a certain layer in a certain view and they determine the effects and behavior of the available layout settings (General, Disconnected, Hierarchical, Orthogonal, Symmetric, and Routing).

The Layer Setup option provides you with default values for each parameter, which you can leave or change as required. The parameter values are in pixels.

This section includes the following topics:

General Layout

The General Layout settings are settings that are available with all layout styles. If your layout contains more than one graph, each graph is laid out separately, and then the intergraph edges (the relationships between the graphs) are routed.

In this tab you define general layout parameters.

➤ **Margin Spacing** – Sets the borders around each node in the layer.

➤ **Constant Spacing** – Sets a constant space around the displayed layer, vertically and horizontally. This space remains the same even if the size of the layer changes.

- **Left** – The size of the space to the left of the display. Default is 20.

- **Top** – The size of the space above the display. Default is 20.

- **Right** – The size of the space to the right of the display. Default is 20.

- **Bottom** – The size of the space under the display. Default is 20.

➤ **Proportional Spacing** – Sets a space around the displayed layer; the space is proportional to the display's size. This space changes when the size of the display changes.

- **Left** – The size of the space to the left of the display. Default is 0.

- **Top** – The size of the space above the display. Default is 0.

- **Right** – The size of the space to the right of the display. Default is 0.

- **Bottom** – The size of the space under the display. Default is 0.

➤ **Nested View Spacing** – For future use.
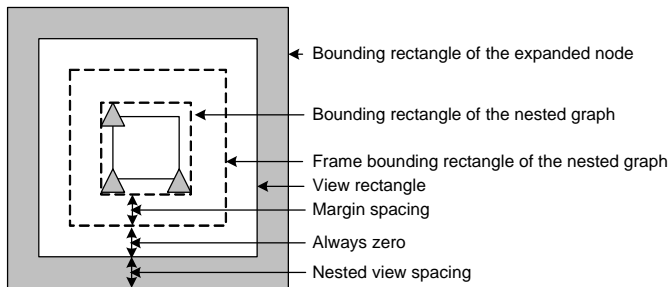
➤ **Intergraph Edge Spacing** – If intergraph edges are routed orthogonally, you can specify horizontal and vertical spacing for the edges.

   ➤ **Horizontal Spacing** – Sets the horizontal distance between edges.

   ➤ **Vertical Spacing** – Sets the vertical distance between edges.

Disconnected Layout

In this tab, you define general layout parameters for disconnected nodes and connected components:

➤ A disconnected node is a node that is not connected to any other node.

➤ A disconnected component is a component that is not connected to any other component. It contains a set of nodes that can be connected to each other.

The following example illustrates six connected components:

A connected component with five nodes

A connected component with two nodes

Four connected components, each made up of a single disconnected node
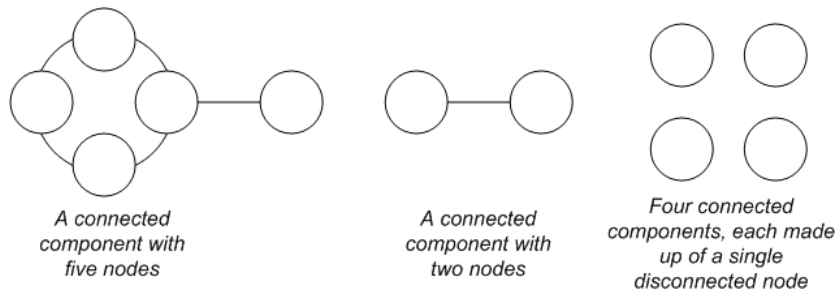
➤ **Components** – The spacing between the components in a disconnected graph consists of both a constant value and a proportional value based on the sizes of the components.

   ➤ **Detect Components** – Select **Detect Components** to view the disconnected components. You can specify that all components be laid out as together or individually, without regards to the other components. If the grouping for components is selected, each component is laid out, and then the resulting components are packed together.

   ➤ **Constant Spacing** – Sets constant spacing (horizontal and vertical) around each disconnected component. Default is 20.
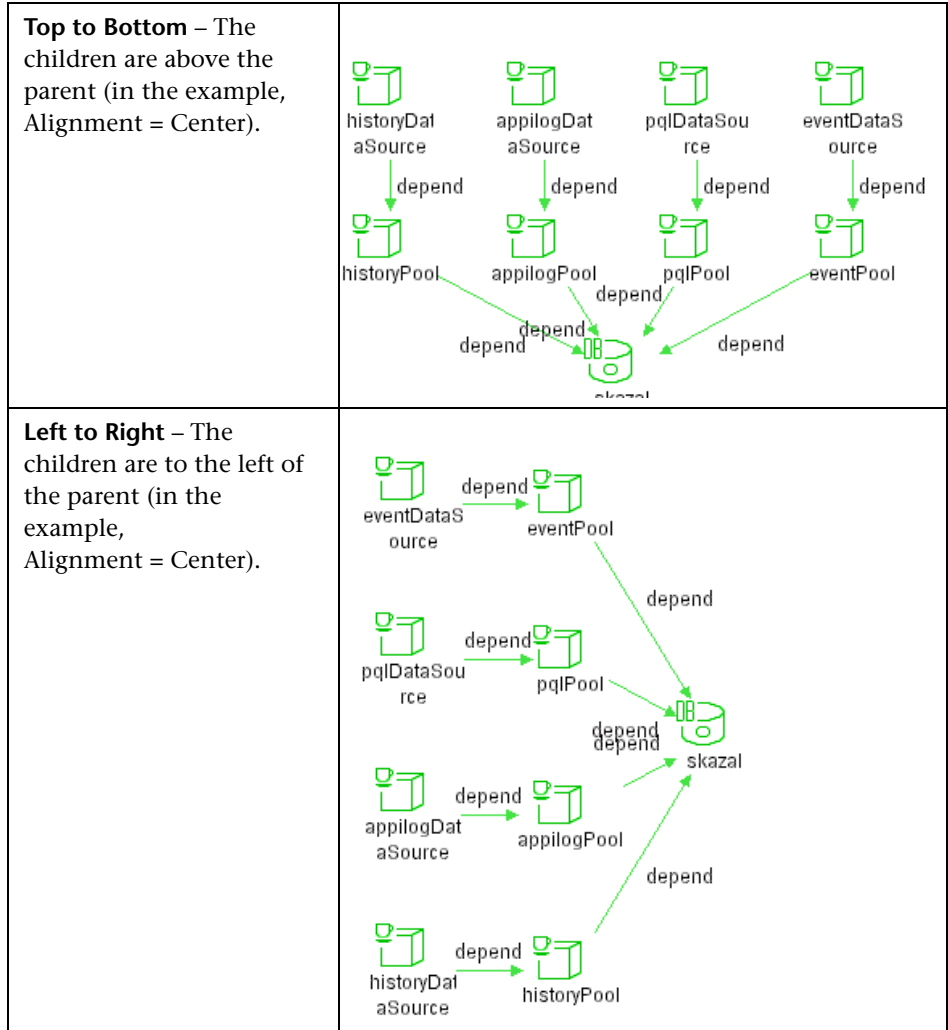
> ➤ **Proportional Spacing** – Sets spacing (horizontal and vertical) that is proportional to the size of the component, around each disconnected component. Default is 0.05.

➤ **Disconnected Nodes** – The spacing between the nodes in a disconnected graph consists of both a constant value and a proportional value based on the sizes of the nodes.

> ➤ **Detect Disconnected Nodes** – Select **Detect Disconnected Nodes** to view the disconnected components. You can specify that disconnected nodes be grouped into one component or treated individually while being laid out.

> ➤ **Constant Spacing** – Sets constant spacing (horizontal and vertical) around each disconnected node. Default is 20.

> ➤ **Proportional Spacing** – Sets spacing (horizontal and vertical) that is proportional to the size of the component, around each disconnected node. Default is 0.05.

➤ **Aspect Ratio** – The layout's proportions. It only affects the placement of components with respect to each other, not the layout of each individual component. It has no effect if the graph is connected.

> ➤ **Automatic** – The layout is performed automatically.

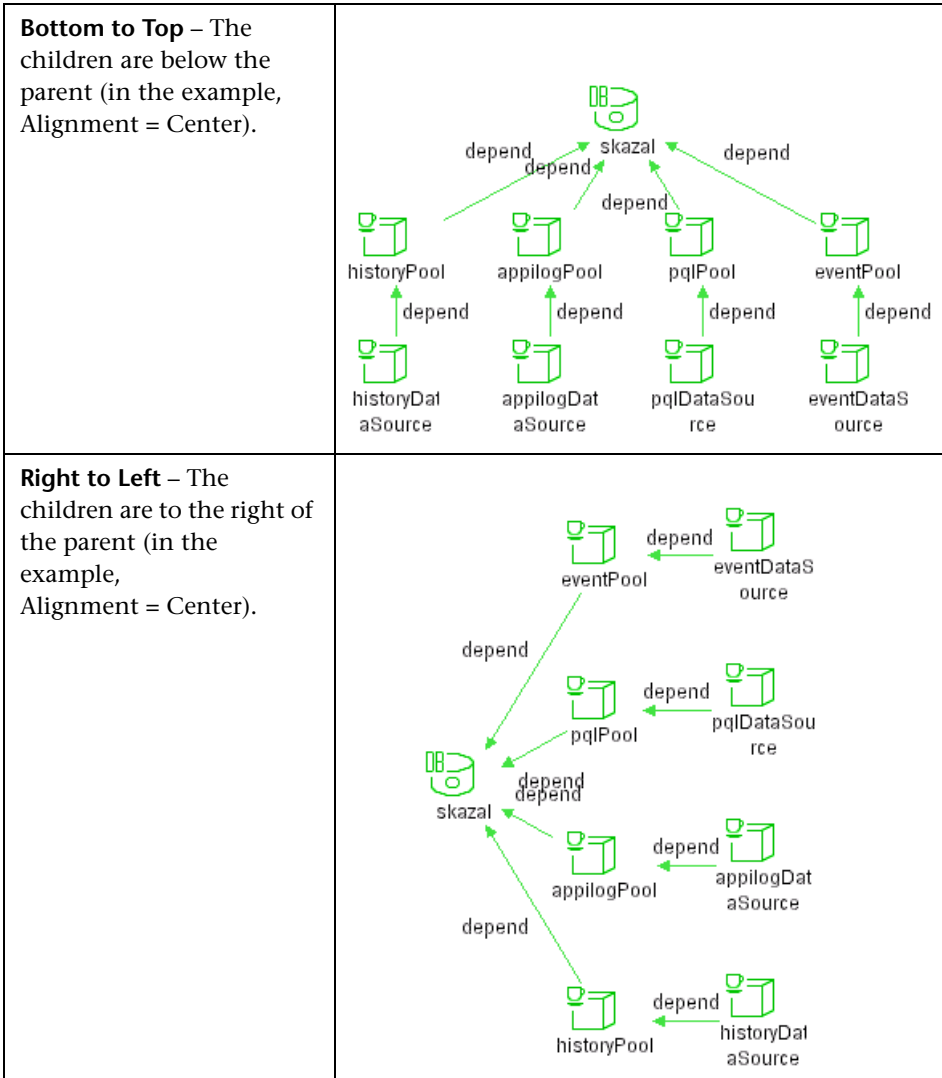> ➤ **Custom** – You can specify the aspect ratio. The default is 1.0.

Hierarchical Layout

The hierarchical layout shows the precedence relationships that can represent organizational or information management system dependencies, as well as process models, software call graphs, and work flows. The hierarchical layout emphasizes dependencies by placing the nodes at different levels.

There are several layout options you can apply to a Hierarchical layout, as described below.
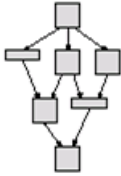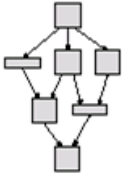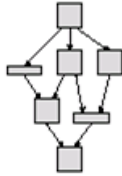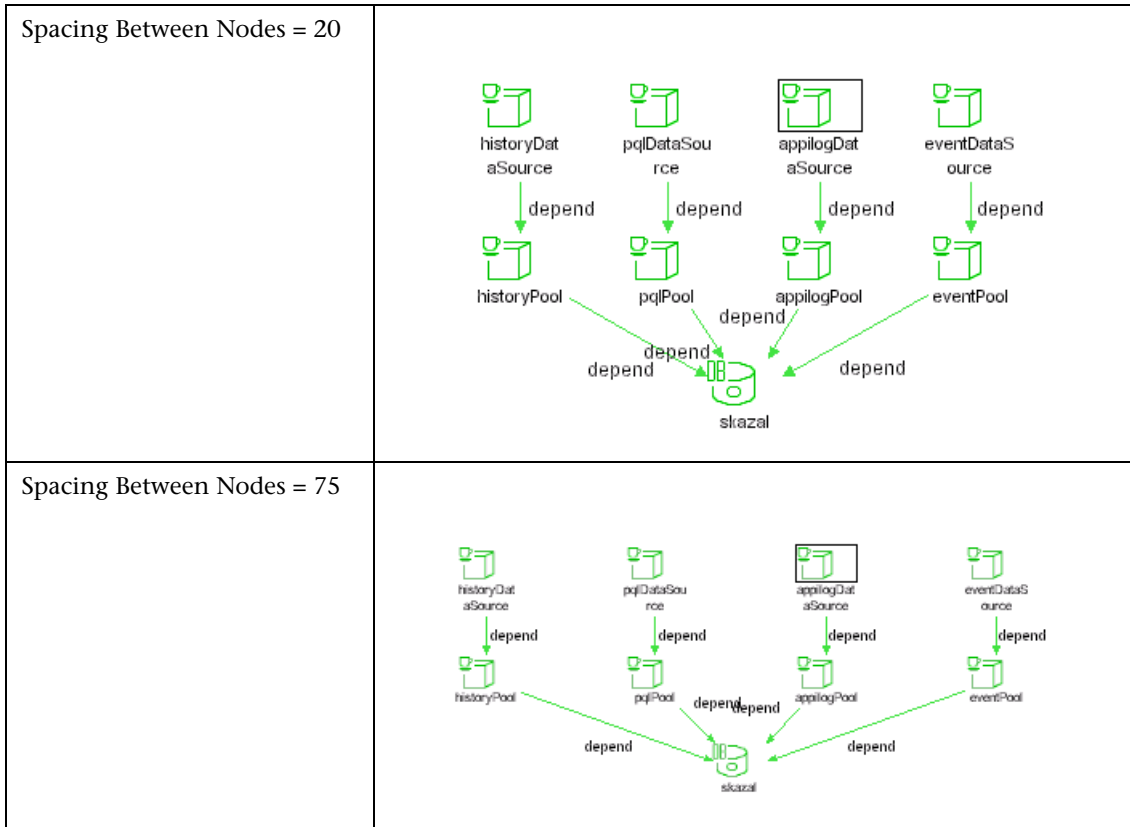
➤ **Orientation** – The orientation of the hierarchy.

Top to Bottom – The children are above the parent (in the example, Alignment = Center).	
Left to Right – The children are to the left of the parent (in the example, Alignment = Center).	

183

Bottom to Top – The children are below the parent (in the example, Alignment = Center).	
Right to Left – The children are to the right of the parent (in the example, Alignment = Center).	

➤ **Level Alignment –** This option allows you to control the way nodes are aligned on a vertical level.

Top – The top of the nodes at the same level of the hierarchy are aligned.	**Center** – The center of the nodes at the same level of the hierarchy are aligned.	**Bottom** – The bottom of the nodes at the same level of the hierarchy are aligned.

➤ **Layout Quality –** You can adjust the quality of the layout produced to fit your application's needs. The quality reflects the number of steps or the method used to produce the layout (for example: high quality is neater and cleaner but takes more time to layout).
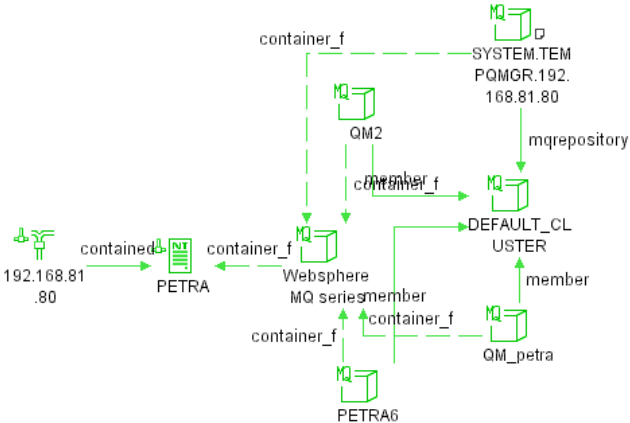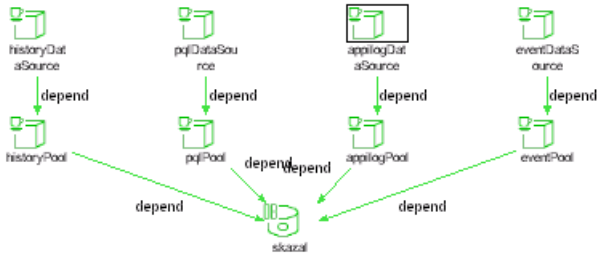
> ➤ **Draft**

> ➤ **Default**

> ➤ **Proof**

> ➤ **Horizontal Spacing –** Sets the minimum horizontal distance between two neighboring nodes on each level.

> ➤ **Spacing Between Nodes** – To set a space around each node–horizontally–in the layer. The default value is 20.
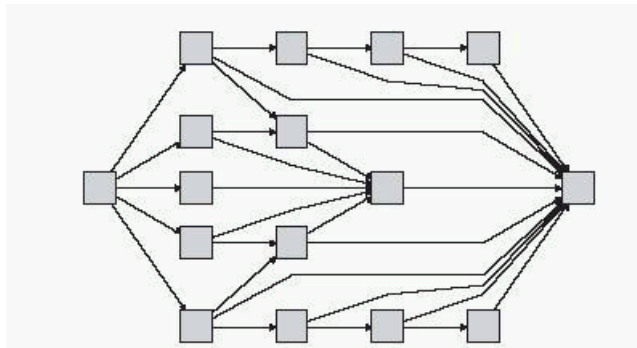
Spacing Between Nodes = 20	
Spacing Between Nodes = 75	

➤ **Spacing Between Edges** – This setting is available only when you select **Routing** > **Orthogonal**. It sets a horizontal space between elements in the graph.

Spacing Between Edges = 20	
Spacing Between Edges = 50	
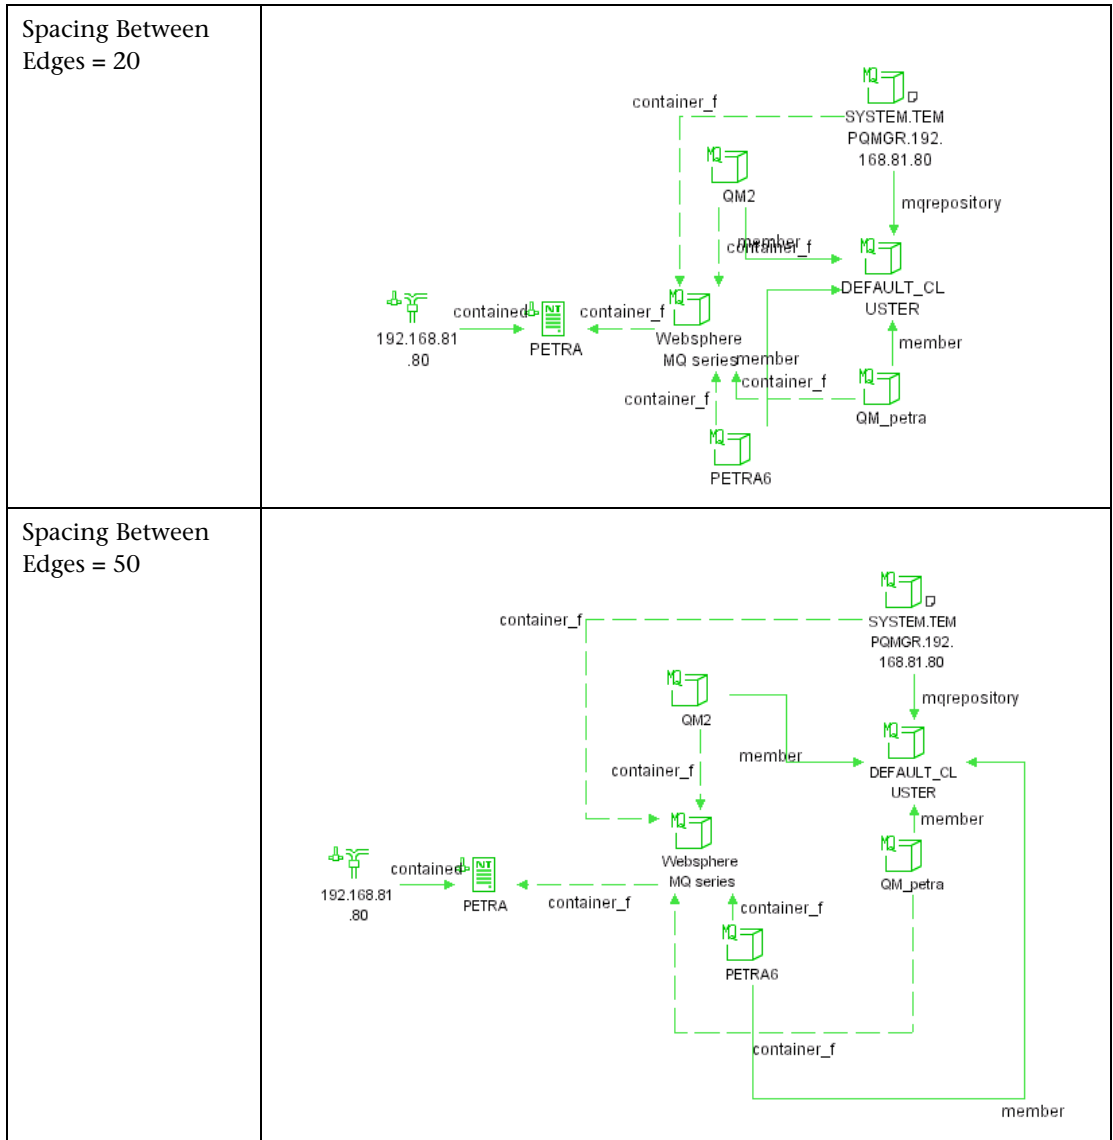
➤ **Vertical Spacing** – Sets the minimum vertical distance between two neighboring nodes on different levels.

➤ **Spacing Between Nodes** – Sets a vertical space around each node in the layer. The default value is 20.

➤ **Spacing Between Edges** – This setting is available only when you select **Routing** > **Orthogonal**. It sets a vertical space between edges in the graph.

➤ **Routing** – To set the type of relationship between CIs. This option is useful when layers have nodes with a very large number of connecting relationships.

Orthogonal – To have orthogonal relationships between CIs.	
Polyline – To have straight lines between CIs.	

➤ **Orthogonal Routing** – Routes edges horizontally and vertically.

➤ **Fix Node Sizes** – For future use.

➤ **Polyline Routing** – Routes edges as one or more straight line segments with arbitrary angles. Polyline routed edges are prevented from overlapping by adding extra segments.

➤ **Spacing Between Bends** – This setting is available only when you select **Routing** > **Polyline**. Polyline routing routes edges as one or more straight line segments with arbitrary angles. Path nodes are added automatically to keep the edges from overlapping.


➤ **Variable Level Spacing** – Adjusts the spacing between pairs of neighboring levels according to the density of edges between the levels.

If the drawing's edges are orthogonally routed, this ensures the required vertical spacing between horizontal edges (in a top-to-bottom or bottom-to-top layout) routed between levels.

If the routing is polyline, variable level spacing makes it easier to distinguish between edges in very dense drawings.

➤ **Undirected Layout** – Edge direction is not used to build the levels of the hierarchical drawing.

Orthogonal Layout

The orthogonal layout routes relationships horizontally and vertically. This results in relationships bending at ninety-degree angles only.

➤ **Horizontal Spacing** – To set the minimum horizontal distance between CIs.
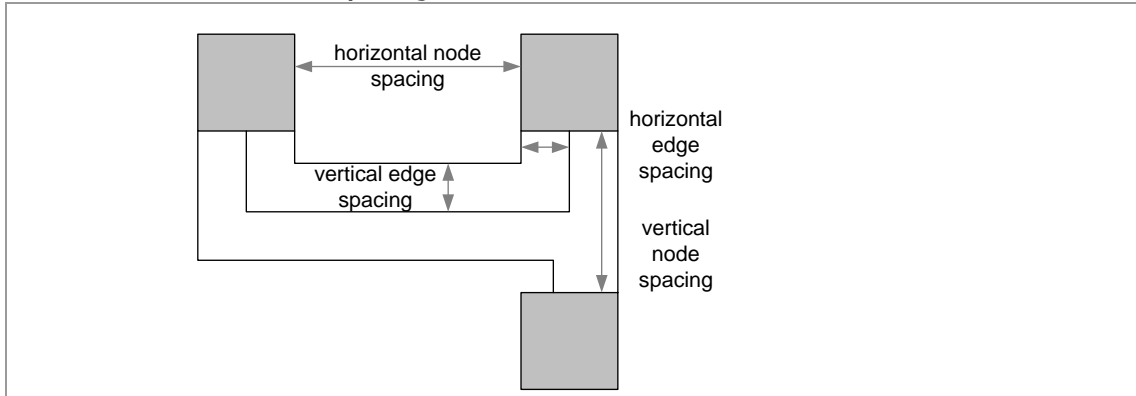
➤ **Spacing Between Nodes** – Sets a space around each node in the layer. The default value is 20.

➤ **Spacing Between Edges** – Sets a space between elements in the layer.

➤ **Vertical Spacing** – To set the minimum vertical distance between CIs.

➤ **Spacing Between Nodes** – Sets a space around each node in the layer. The default value is 20.

➤ **Spacing Between Edges** – Sets a space between elements in the layer.

➤ **Layout Quality** – You can adjust the quality of the layout produced to fit your application's needs. The quality reflects the number of steps or the method used to produce the layout (for example: high quality is neater and cleaner but takes more time to layout).

➤ **Draft**

➤ **Default**

➤ **Proof**

➤ **Aspect Ratio** – The layout's proportions. It only affects the placement of components with respect to each other, not the layout of each individual component. It has no effect if the graph is connected.

➤ **Automatic** – The layout is performed automatically.

➤ **Custom** – Specify the aspect ratio. The default is 1.0.

➤ **Disabled** – To disable the aspect ratio.

➤ **Fix Node Sizes –** For future use.

Symmetric Layout

The symmetric layout emphasizes the symmetries that may occur in a graph.
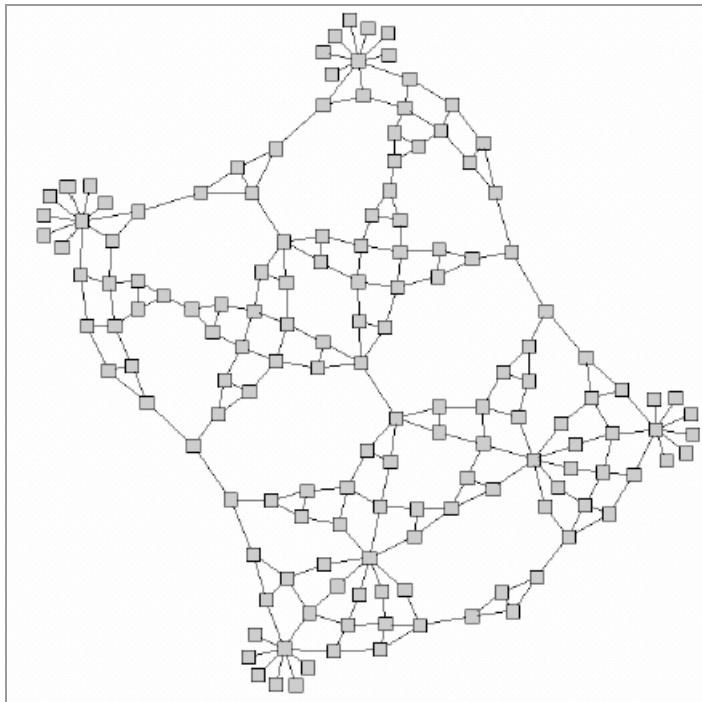
This section describes the symmetric layouts that are available.

➤ **Spacing Options** – The symmetric layout allows you to adjust the spacing of nodes.

➤ **Node Spacing** – Allows you to set constant horizontal and vertical spacing around each node in the layer. The value you specify is a guideline for the layout, so that it is possible that the spacing for a particular pair of nodes might be different from the one you specified. The larger the node spacing you specify, the farther apart the nodes will be in the final layout. The default is 10.
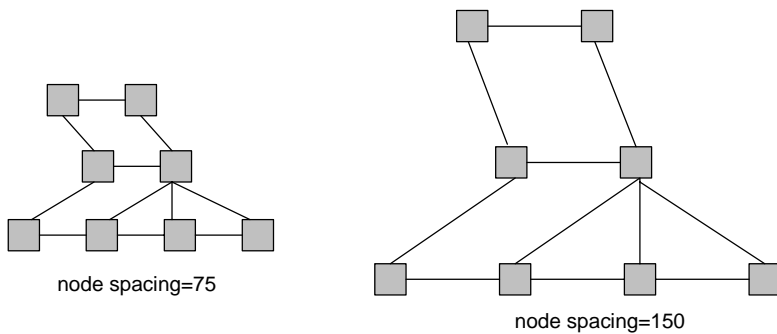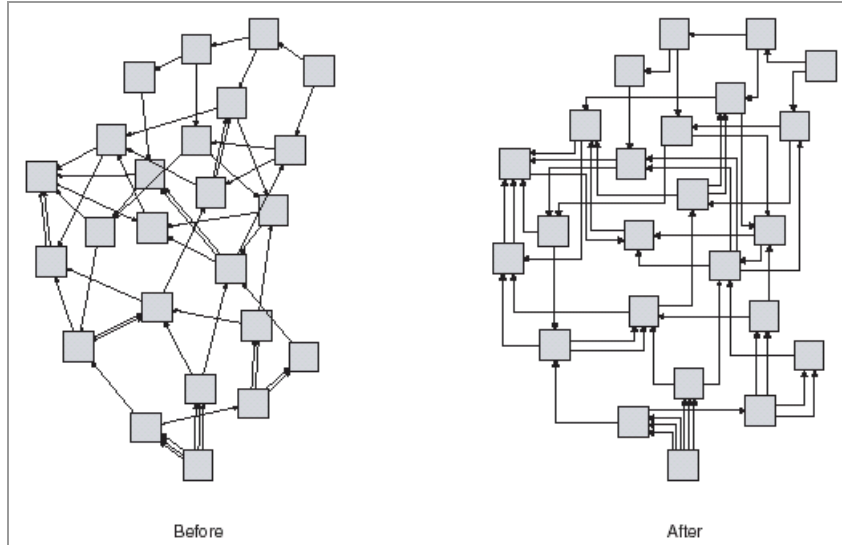

node spacing=75

node spacing=150

➤ **Layout Quality** – You can adjust the quality of the layout produced to fit your application's needs. The quality reflects the number of steps or the method used to produce the layout (for example: high quality is neater and cleaner but takes more time to layout).

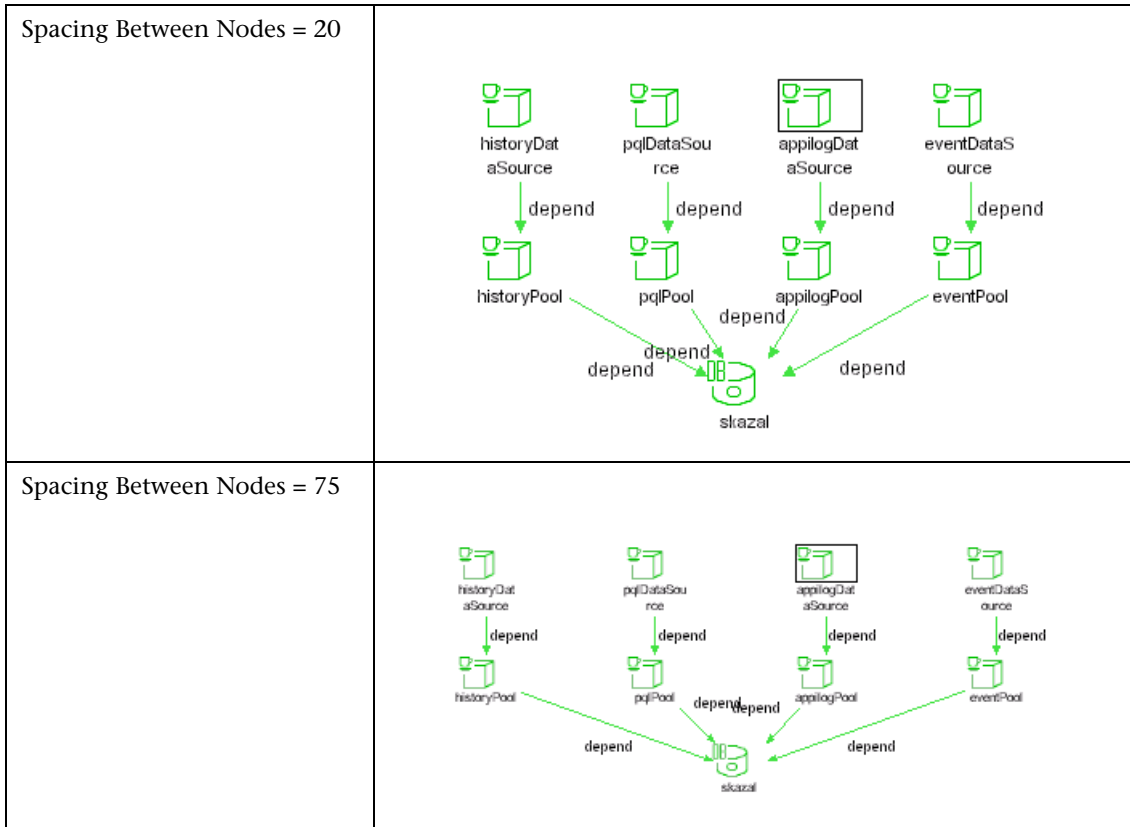➤ **Draft**

➤ **Default**

➤ **Proof**

Routing

Routing produces drawings that leave nodes essentially where they are and reroutes the edges orthogonally.


Before   After

This section describes the routing layouts that are available.

> ➤ **Horizontal Spacing** – To set the horizontal space around each node or each CI.

>> ➤ **Spacing Between Nodes** – Sets the space around each node–horizontally–in the layer. The default value is 20.

Spacing Between Nodes = 20	
Spacing Between Nodes = 75	

> ➤ **Spacing Between Edges** – Sets the horizontal space between any two neighboring parallel edges. Edge spacing can also affect the size of a node to which many edges are attached, unless the node is set to remain fixed at its current size. The settings for edge spacing are constant values.

For example:

Spacing Between Edges = 20	
Spacing Between Edges = 50	
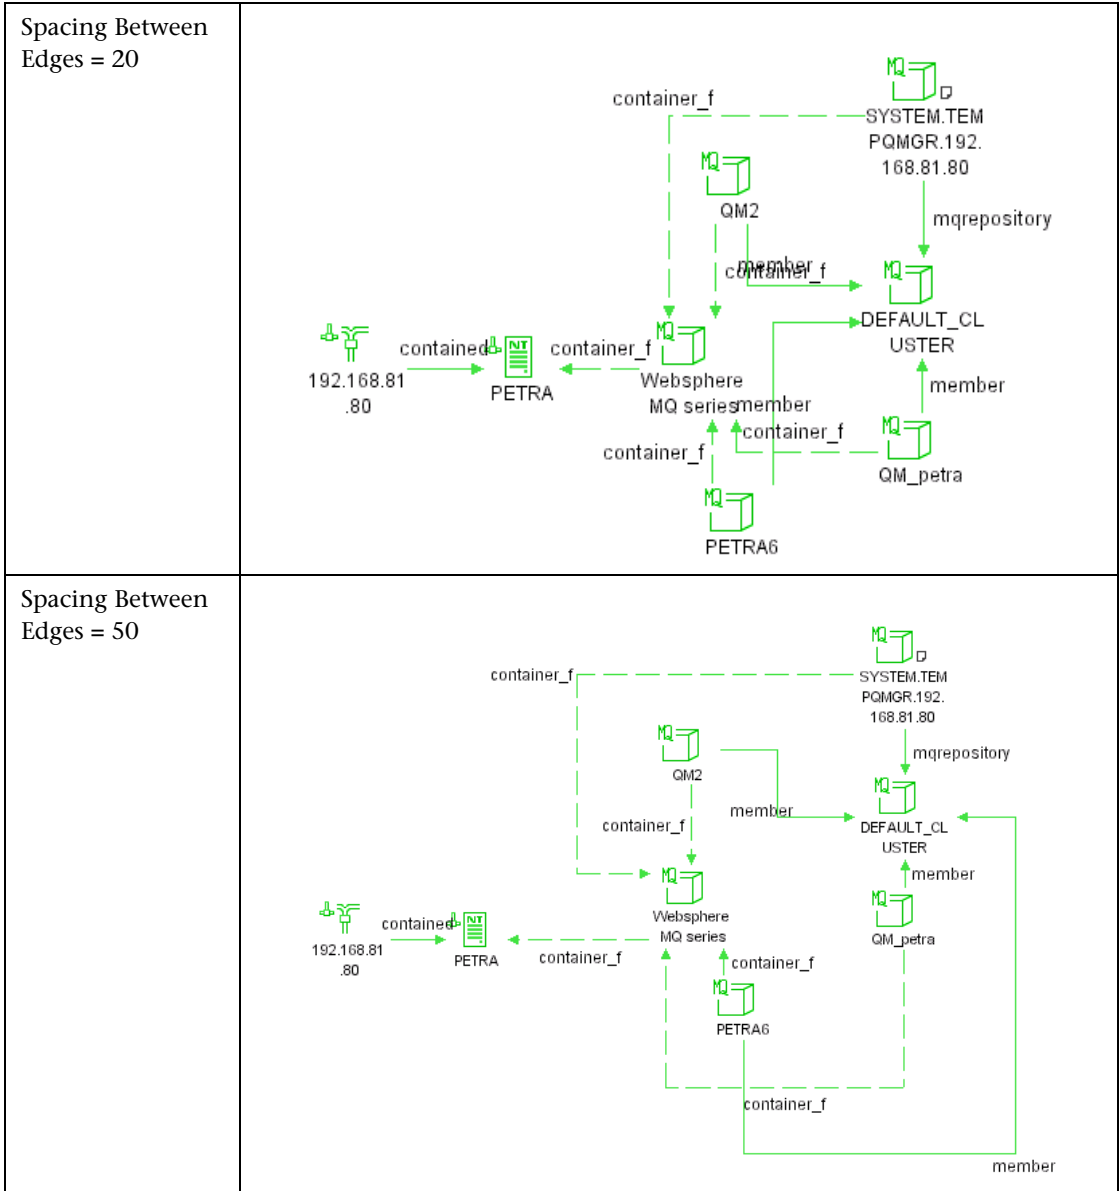
➤ **Vertical Spacing –** To set the vertical space around each CI.

➤ **Spacing Between Nodes** – Sets the space around each node–vertically–in the layer. The default value is 20.

➤ **Spacing Between Edges** – Sets the vertical space between any two neighboring parallel edges. Edge spacing can also affect the size of a node to which many edges are attached, unless the node is set to remain fixed at its current size. The settings for edge spacing are constant values.

➤ **Fix Node Sizes –** Select **Fix Node Sizes** to maintain the size of a node. Clear the option if you want the layout to increase the size of a node if it is necessary to maintain the specified relationship spacing when more than one relationship is attached to the same node side.

➤ **Fix Node Positions –** Select **Fix Node Positions** to maintain the position of all the nodes in the display. Clear the option if you want the layout to move the nodes if it is necessary to improve the layout. The movement is minimal to avoid overlaps and minimize bend points.

Deactivating and Activating Views

You can deactivate a view if you do not currently need it in Mercury Application Mapping. For example, you can deactivate a view if the view was created for test reasons and it is not being accessed by any user. When a view is deactivated, the view is not built in the CMDB and as a result, the load on the system is reduced.

Views are automatically deactivated if the view contains more than the maximum number of CIs allowed by CMDB thresholds. Views are also deactivated if there are more than 200 active views, the default number of views. In the latter case, you can reactivate the current view if you delete or deactivate other unwanted views to reduce the number of active views to less than 200.

To deactivate or activate a view:

1 In the Topology Map or Information pane, right-click the selected view and select **Deactivate** to deactivate an active view. Select **Activate** to activate an inactive view..

2 When a view is deactivated, all CIs associated with the view disappear. When a view is activated, all CIs associated with the view appear. There may be a slight delay until the pane updates the views.

Improving Views

You can improve a view.

➤ Modifying the default layer setup to achieve the clearest presentation possible

➤ Resizing the labels of all CIs that are displayed in one layer in the topology map

➤ Adding an image file (of .gif type) as a background image to be used as a navigational aid

➤ Adding notes to nodes or events in the topology map.

This section includes the following topics:

Defining a Layer Layout

You can customize the layer layouts by modifying the values of the default layer layouts.

To define a layer layout:

1 From the **Map** menu, select **Layer Setup** to open the Layer Properties dialog box with the following tabs:

➤ **General** – Defines general layout types for all layout types. For details, see "General Layout" on page 180.

➤ **Disconnected** – Defines general layout parameters for disconnected nodes and connected components. For details, see "Disconnected Layout" on page 181.

➤ **Hierarchical** – Shows the precedence relationships of the nodes in the selected layer. For details, see "Hierarchical Layout" on page 182.

➤ **Orthogonal** – Routes relationships only horizontally and vertically. For details, see "Orthogonal Layout" on page 190.

➤ **Symmetric** – Displays a clear representation of complex networks. For details, see "Symmetric Layout" on page 191.

➤ **Routing** – Leave nodes where they are in the graphic representation and reroutes the edges orthogonally. For details, see "Routing" on page 193.

2 Enter new values for the layout type you want to change. See the following sections for details on each layout type.

➤ Click **Defaults** to restore the original default values.

➤ Click **Apply** to apply the changes you have made. The changes you made are saved but the dialog box remains open.

➤ Click **OK** to save the changes you have made to the database and close the dialog box.

Adding a Background Image to a Topology Map

You can add an image file (of .gif type) as the background image of a view. This can be used as a navigational aid, for example, for global networks containing different networks in different countries.

To add a background image:

1 In the View Explorer pane, select the required view and then select **Manual Layout** from the list in the toolbar or click **View** > **Layout** > **Manual** to display the manual layout view option.

2 From the **Edit** menu, select **Add Background Image** to open the Add Background Image dialog box.

3 From the **File Name** area, select the image that you want to display in the background of the selected view.

4 Click **OK**. The selected background image is displayed in the topology map.

To remove a background image:

Display the Add Background Image dialog box, and select **(none)** from the File Name area.

To add background images to the existing list:

Insert images (of .gif type) in the following location:
Mercury\MAM\root\lib\gui\images\classes\backgrounds

Adding Notes to Displayed CIs

You can create personal notes that include comments, and attach them to nodes or events in the topology map. You can use these notes as reminders to yourself or to provide feedback to other users. You can edit and delete existing notes.

Adding a Note

You can add a note to a CI.

To add a note to a CI:

1 From the Explorer, Map, or Information pane, right-click the CI to which you want to add a note and select **Note** > **Add Note** to open the Note dialog box.

2 Type the note text in the dialog box.

3 Click outside the Note dialog box or click **OK**.

Displaying and Editing a Note

CIs with notes have an indicator in their bottom right corners:

SBLAPP1

To display and edit a note:

1 Double-click the indicator in the bottom right corner of the CI or right-click the CI and select **Note** > **Add Note** to open the Note dialog box.

2 To edit the note's text, type your changes in the dialog box and click **OK**.

Deleting a Note

Right-click the CI and select **Note** > **Delete Note**.

Viewing and Editing CI Attributes

The Attributes dialog box enables you to view and edit specific CI attributes, depending on your user access rights to the system (for details, see Chapter 1, "User and Role Management,"in the *Mercury Application Mapping Administration Guide*).
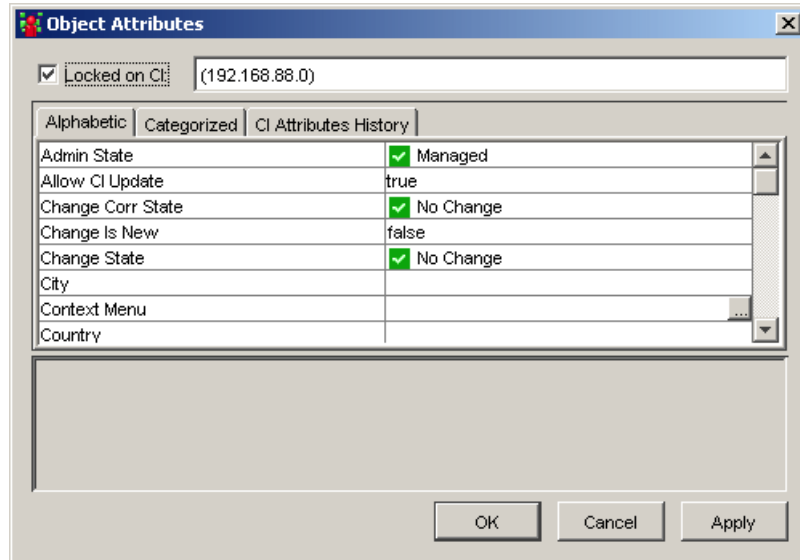
The CI Attributes dialog box enables you to view and edit the attributes defined for the selected CI. The name of the selected CI appears at the top. When the **Locked on CI** check box is checked, Mercury Application Mapping displays only the attributes of the CI in the box at the top. Clear the **Locked on CI** check box to display the attributes of any CI you select.

Note: Only the attributes that were marked as **Editable** and **Visible** in the CI Type Manager appear in the CI Attributes dialog box. For details, see "Creating CITs" on page 384.

The CI Attributes dialog box has the following tabs:

➤ **Alphabetic** – Displays the CI's attributes in alphabetical order.

➤ **Categorized** – Displays the CI's attributes according to the CIT to which it belongs. Click the **Expand** button to display the CI attributes for that CIT.

➤ **CI Attributes History** – Displays a list of attributes in which changes have occurred.

Note: The CI Attributes History tab only displays the attributes whose values have been defined as **Change Monitored** in the Add Attributes dialog box in the CI Type Manager. For details, see "Creating CITs" on page 384.

The CI Attributes History tab has the following fields:

Field	Description
Attribute	The name of the attribute.
Change Date	The date when the change occurred.
Changer	Indicates the cause of the change. For example, the field can contain the name of a user or a discovery pattern.
New Value	The new value of the attribute.

Note: You can reduce the number of attributes that appear in the CIs Attributes History tab by selecting an attribute from the **Filter by** list.

To view and edit CI attributes:

1 Edit the attributes as required in either the **Alphabetic** or **Categorized** tabs. You can only edit the attribute if a pencil icon appears at the bottom next to the attribute name. To edit the value of an attribute, click the value cell and then click the square button on the right.

2 Make the required changes in the dialog box that opens and click **OK**. The changes appear immediately.

3 Click **Apply** in the CI Attributes dialog box to save your changes or click **OK** to save your changes and close the CI Attributes dialog box.

Viewing Layers with a Large Number of CIs

A layer cannot display more than 1000 CIs in one topology map. You can, however, display layers that contain more than 1000 CIs and less than 3001 in a table format. When using this option, the layer's CIs are displayed in a Details window, and can be viewed in successive segments (for using the Details window, see the previous section).

If the layer contains more than 3000 CIs, it cannot be opened. If you try to open it, a message is displayed. Click **OK** to confirm the message. In this case, it is recommended to use organization rules to spread the CIs among several layers. For using organization rules, see the previous section.

To display a layer that has more than 1000 CIs:

1 Double-click the layer you want to open. A message is displayed that the layer cannot be displayed as a map. Click **OK** to display the layer in the Details window.

2 Click **OK.** The Details window is opened, displaying the first 1000 CIs of the layer:

The arrows in the lower-right corner of the Details window enable you to display the other CIs that are in the layer in quantities of 1000 CIs.

Note: You can configure the maximum number of CIs that can be displayed in a table format (default is 3000), and the number of CIs that are displayed in each bulk in the Details window (default is 1000). For details on configuring these parameters, see Map Server Parameters in *Mercury Application Mapping Administration Guide*.

Displaying Interdependent CIs

The **Get Related CIs** option enables you to display the interdependencies of a selected CI in different contexts: a specific layer, a specific view or the entire CMDB. By using this option, all CIs that are connected to the selected CI in a chosen context are displayed, either on the topology map or in a separate window. This way, you can visualize, locate and manage the interdependencies that exist among your IT resources.

The **Get Related CIs** option opens a submenu, which contains three main commands—one for each interdependency context—Layer, View, and Database.

Selecting one of the **Get Related CIs** commands opens an interdependency map for the selected CI. Subsequently, Mercury Application Mapping provides you with additional commands, which enable you to expand the interdependency data that the map displays to other connected CIs, and get a broader perspective on the structure and state of your infrastructure.

Depending on the display option you selected, an additional command appears on the shortcut menu: **Next on Layer**, **View**, or **Database**. These commands are enabled only after using the **Layer**, **View**, or **Database** options.

By applying one of these commands to a connected CI, additional interdependency information is displayed. You can then continue using this command, and display interdependency data for any displayed CI.
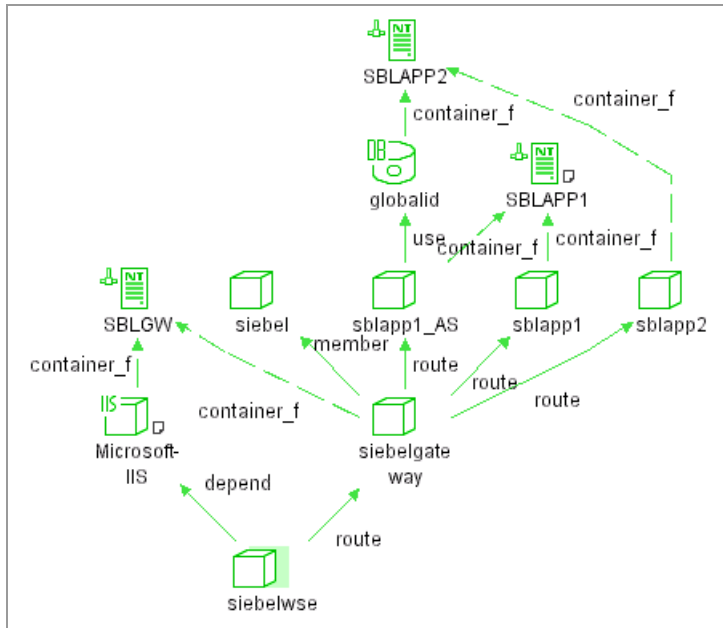
To display interdependent CIs:

1 In the topology map, View Explorer or Information pane, right-click the CI whose interdependent CIs you want to display.

2 Click **Get Related CIs** and then select one of the following commands:

➤ **Layer** – Displays all CIs that are connected to the selected CI via relationships in a specific layer. This option actually removes from the displayed layer all CIs that are not connected to the selected CI, thus enables you to focus on the interdependencies of a specific CI on a given layer.

Note: Once you select the **Layer** option, the new display remains as the layer's default until the view is rebuilt. To restore the original display, click the **Show Hidden Symbols** button on the toolbar.

The following example shows the way the siebelgateway CI is displayed on a regular view's layer:

When using the **Layer** option, Mercury Application Mapping hides all CIs that are not connected to siebelgateway on this layer, and displays only its interdependent CIs:
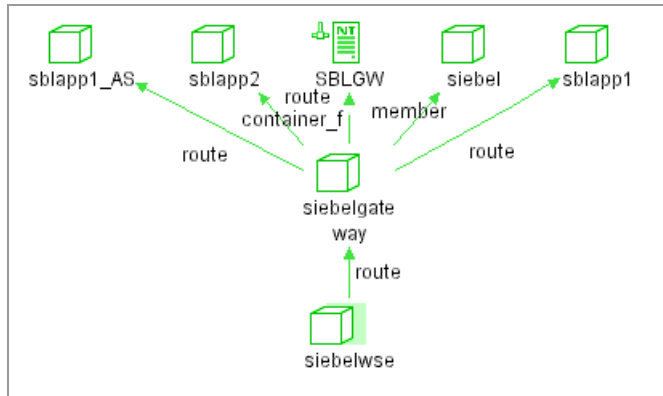

➤ **View** – Displays all CIs, connected to a selected CI, that are part of the view's TQL, regardless of the layer in which they are located according to the view's organizational rules. The View's content is displayed in a new window.

Note: The window's display is static and cannot be refreshed. If you want to update the display, you need to close the window and reopen it.

For example, when using the **View** option to display all CIs that are connected to router1 in a given view, the following window is displayed:
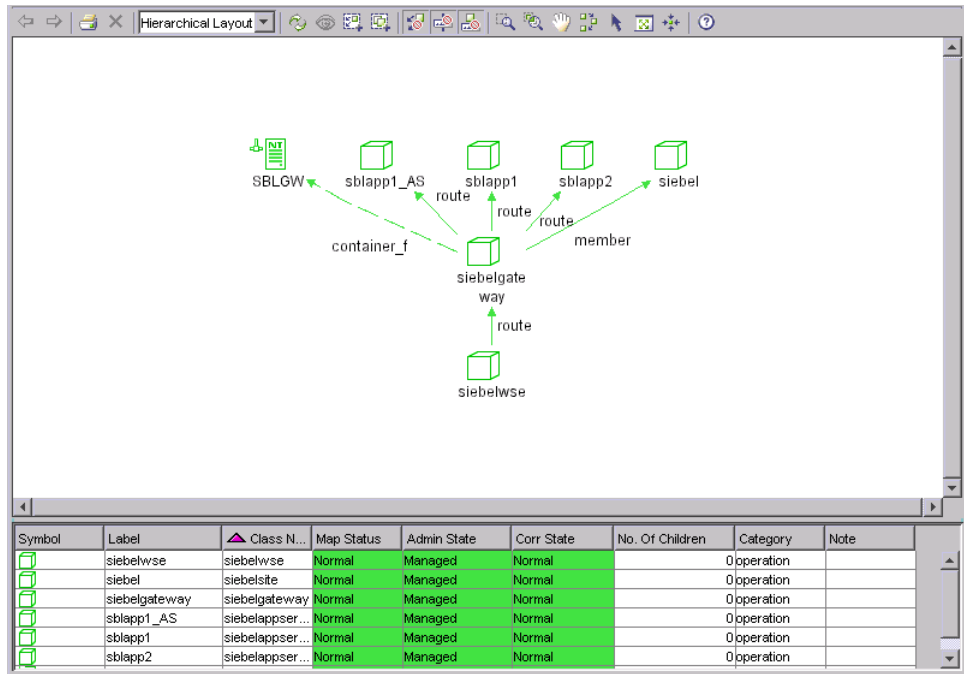

As you can see in the above window, there is one additional CI connected to siebelgateway, in comparison to the Layer display. This CI is located in a lower layer level, and therefore it does not appear on the Layer display.

➤ **Database –** Displays all CIs connected to the selected CI from the CMDB. Similar to the **View** option, the **Database**'s content is displayed in a new window.

The following example shows how, when using the **Database** option, Mercury Application Mapping displays all known CIs that are connected to siebelgateway, regardless of their position in the existing TQL scenarios:
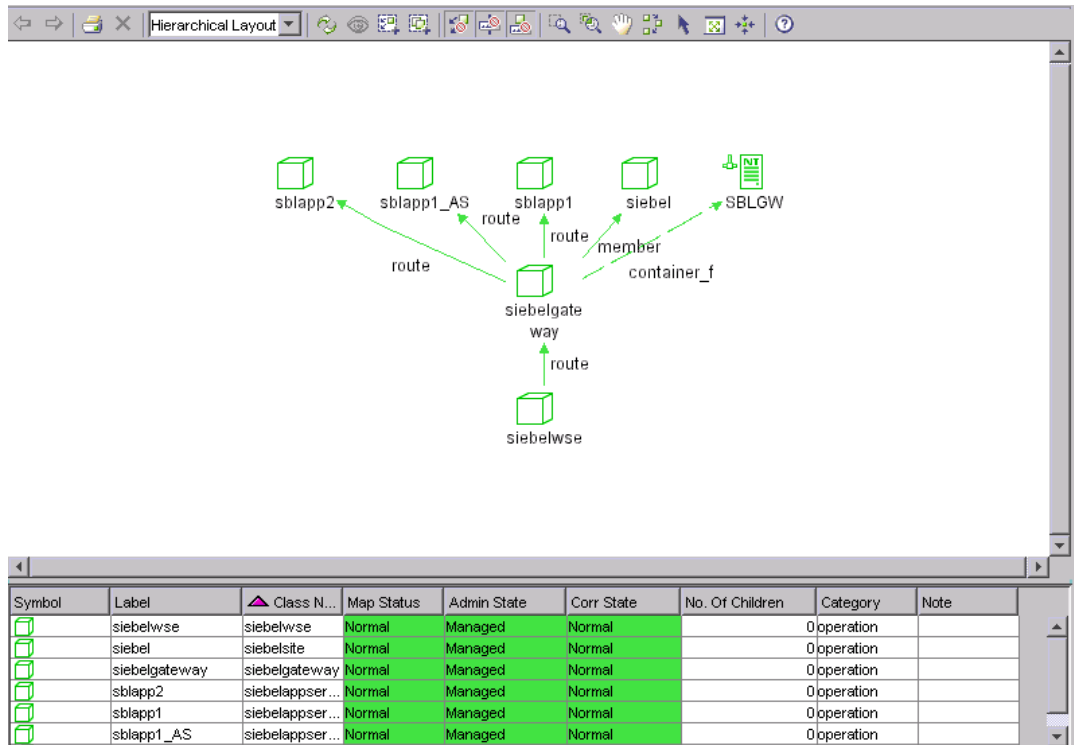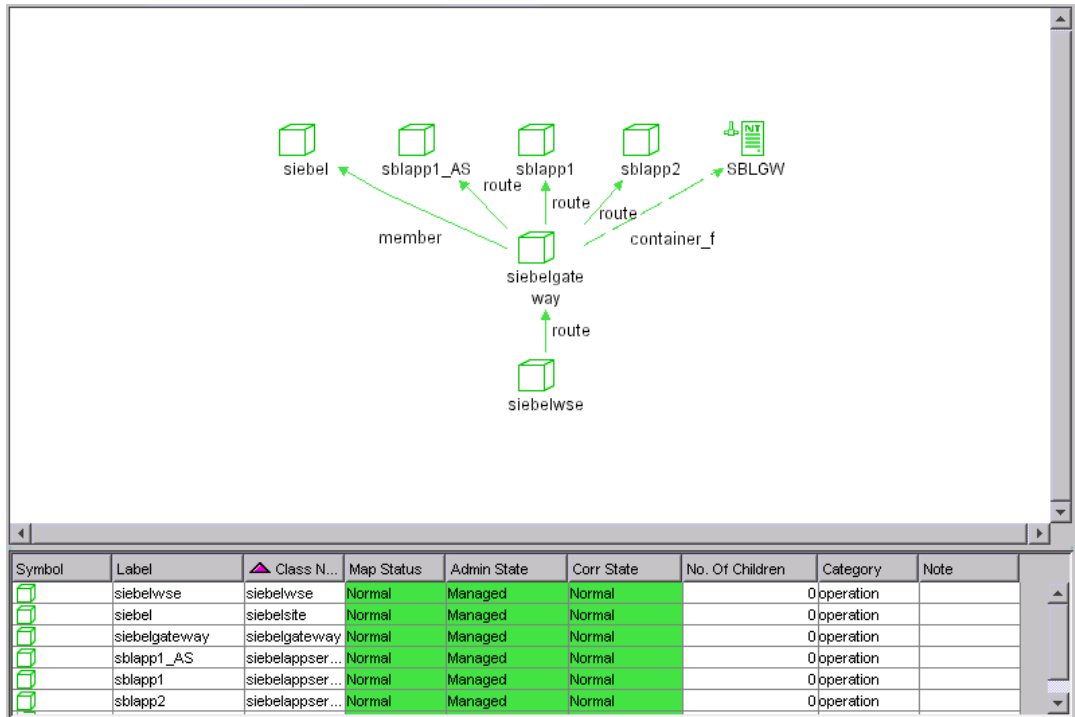

3 (Optional) To add to the displayed map interdependency information about one of the connected CIs, right-click the CI whose interdependencies you want to display. Select **Get Related CIs** and one of the following commands:

➤ **Next on Layer –** (Enabled only after using the **Layer** option) Keeps the previous Layer display on the topology map, and adds to that layer the connected CIs of one of the displayed CIs.

The following example shows how the Next on Layer command is applied to the SBLGW CI, and consequently, its connected CIs are added to the map:
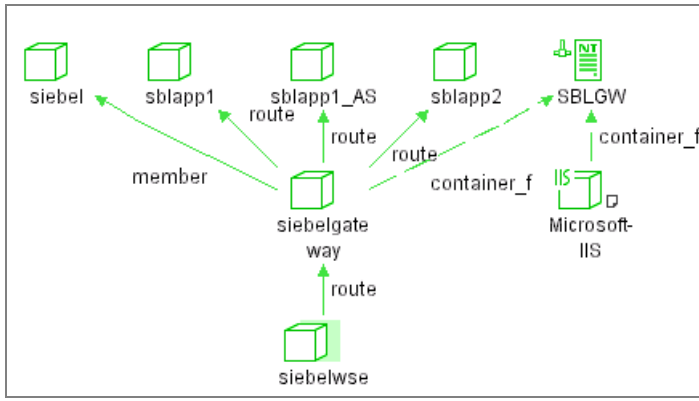
➤ **Next in View –** (Enabled only after using the **View** option) Keeps the previous View display on the topology map and adds the connected CIs from the given view of the selected CI.

Symbol	Label	Class N...	Map Status	Admin State	Corr State	No. Of Children	Category	Note
	siebelwse	siebelwse	Normal	Managed	Normal	0	operation	
	siebel	siebelsite	Normal	Managed	Normal	0	operation	
	siebelgateway	siebelgateway	Normal	Managed	Normal	0	operation	
	sblapp1_AS	siebelappser...	Normal	Managed	Normal	0	operation	
	sblapp1	siebelappser...	Normal	Managed	Normal	0	operation	
	sblapp2	siebelappser...	Normal	Managed	Normal	0	operation	

➤ **Next in Database –** (Enabled only after using the **Database** option) Keeps on the topology map the previous **Database** display and adds to it the connected CIs from the database of the selected CI.
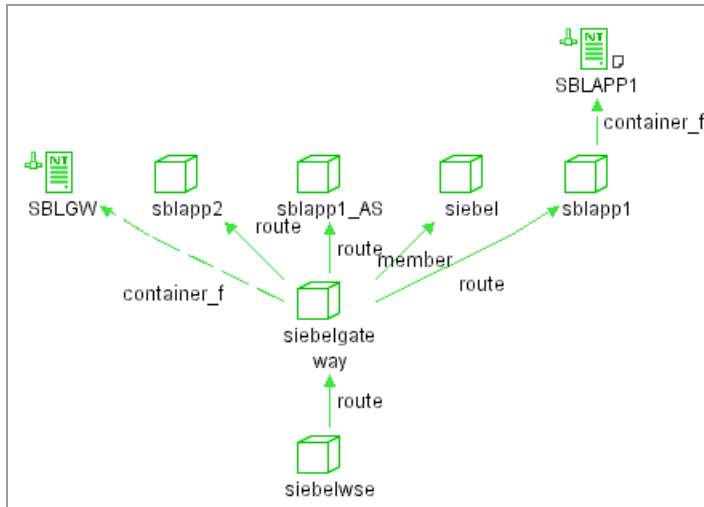
Showing a Correlation Impact

The **Show Impact** option enables you to select a CI, which is defined by a correlation rule as a root cause CI, and to display all the CIs that are affected by it in a separate window.

For example, suppose a correlation rule is defined for an IP, which specifies that the rule affects hosts, ports, and client-servers that are linked to the IP. By using the **Show Impact** option, you can display all CIs that are affected by the IP CI.

Note: For details on defining correlation rules, see Part VI, "Correlation Rules."


To show the correlation impact of a root cause CI:

1 In the topology map or Information pane of Topology Map, right-click a root cause CI and select **Correlation** > **Show Impact**.

➤ If one correlation rule is defined for the selected CI, the Impact window is displayed. Skip to step 2.

➤ If several correlation rules are defined for this CI, the Impact Correlation Rules dialog box is displayed.

The Impact Correlation Rules dialog box lists all the correlation rules that define this CI as a root cause CI.

Select a correlation rule and click **OK**.

2 The Impact window is displayed:
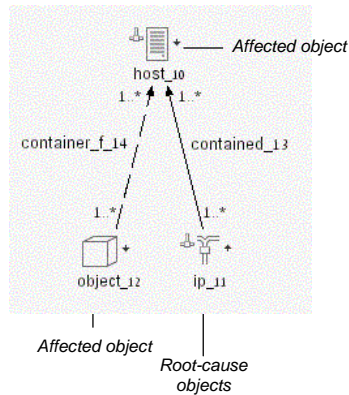
In the above example:

➤ The CPU CI is the root cause CI, and for that reason it is marked by this indicator.

➤ The HAVA CI is part of the correlation rule, but it is neither the root cause CI nor an affected one. For that reason it is marked by this indicator.

➤ The IP CI is the affected CI, and for that reason it is marked by this indicator.

Showing Root Cause CIs

Mercury Application Mapping enables you to select a CI that is affected by a correlation rule, and display the root cause CI(s) that affected it.

This section includes the following topics:

➤ "Showing the Root Cause CI of an Affected CI" on page 214

➤ "Problem Isolation" on page 215

Showing the Root Cause CI of an Affected CI

You can display the root cause CI of an affected CI.

To show the root cause CI of an affected CI:

1 In the Topology Map or Information pane, right-click an affected CI and select **Correlation** > **Show Root cause**.

➤ If one correlation rule is defined for the selected CI, the Correlation window is displayed. Skip to step 2.

➤ If several correlation rules are defined for this CI, the Root Cause Correlation Rules dialog box is displayed.

The Root Cause Correlation Rules dialog box lists all the correlation rules that define this CI as a root cause CI.

Select a correlation rule and click **OK**.

2 The Correlation window is displayed:
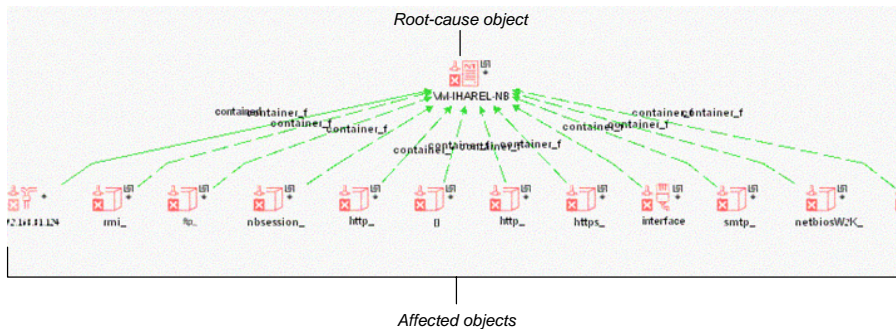
In the above example, the scenario that led to a correlation event occurring for the Switch CI is displayed:

➤ The IP CI on the left is the root cause CI, and for that reason it is marked by this indicator.

➤ The CIs at the bottom are part of the correlation rule, but they are neither the root cause CIs nor the affected ones. For that reason they are marked by this indicator.

➤ The Switch CI at the top is the affected CI, and for that reason it is marked by this indicator.
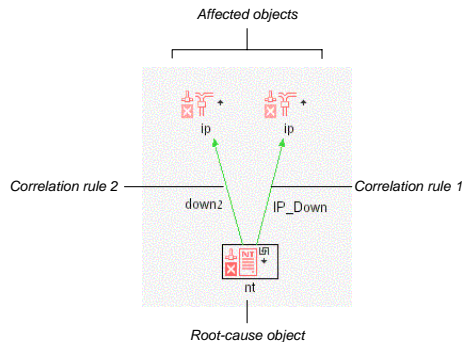
Problem Isolation

This option enables you to retrieve root cause information for CIs that are affected by a chain of correlation rules. The output of problem isolation displays a logical map that describes a chain of trigger/affected CITs that affect the CI.

Show Problem Isolation, as opposed to **Show Root Cause**, allows you to view the source root cause CI in one step. When using the **Show Root Cause** option to reach the root cause CI, you have to use **Show Root Cause** on each trigger/affected rule, one by one.

To show the root cause CI:

In the Map or Information pane, right-click an affected CI and select **Correlation > Show Problem Isolation**. A window, as illustrated below, is displayed containing the root cause CI and all the other CITs in the correlation rule chain.

Use the **Show Problem Isolation** option on the IP CI to view, in one step, that the source of the failure is the Network CI.

Use the **Show Root Cause** option to see that the host CI is the failure as it is the next CIT in the correlation rule chain. To display the Network CI as the root cause CI, you have to right-click the host CI and select **Show Root Cause** again.

To view the information about specific CIs that are triggered or affected by the correlation rule, double-click the relationship connecting the CIs. A map containing the same information as when you use the Show Root Cause option is displayed (for details, see "Showing Root Cause CIs" on page 214).

Using Meta Relationships

A Meta relationship is a logical relationship that represents several physical relationships between CIs that are displayed on different layers. These relationships are not displayed together on one layer, due to the hierarchal structure of the view that organizes CIs and their interdependencies in separate layers. The Meta relationship option enables you to display in one window, the full path of physical relationships that sustain relationships at upper levels.

For example, in a certain view, a Client-Server relationship connects two NTs. However, the actual relationship exists between two ports. These connected ports are not seen from the upper level layer, since they are located by organizational rules in a lower layer level.
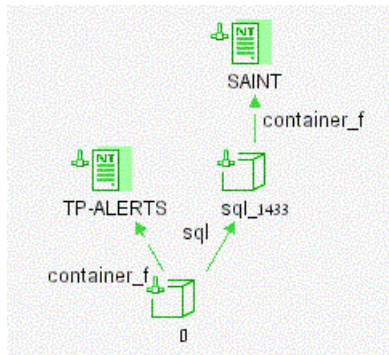
To open and display a Meta relationship:

➤ On the topology map, double-click the Meta relationship of the contents you want to display. A new Relationship Map window appears, displaying the relationships that exist between the two Meta relationship's CIs.

For example the following Meta relationship displays a logical route relationship between a host and a server:

Double-click a Meta Link to open it
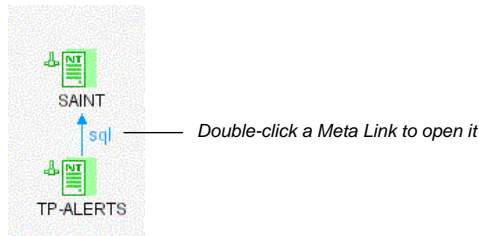
Double-click the route relationship to display all the CIs that compose the physical path:

Note: If you double-click a relationship and a Relationship Map window is not displayed, it means that this relationship is not a Meta Relationship but a regular one.
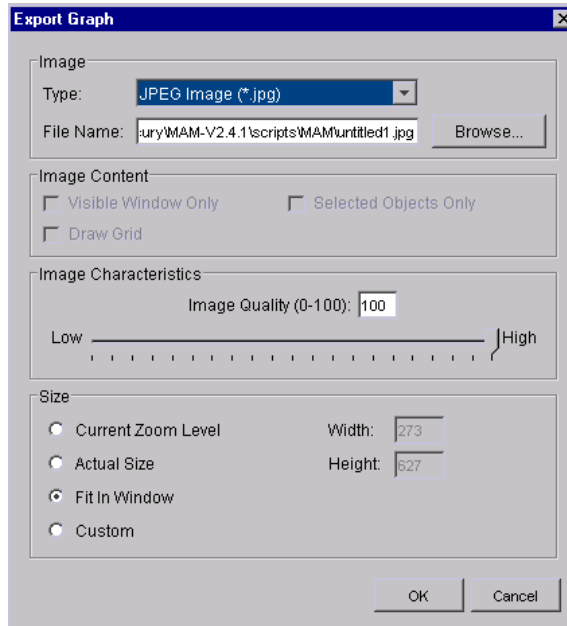

Saving a Topology Map

You can save a topology map to a file.

To save a Topology Map to a file:

1 Select and display the topology map you want to save to file.

**2** Select **Map** > **Export Graph** to open the Export Graph dialog box.

3 In the **Type** list, select the required file format.

4 In the **File Name** box, specify the required file name and location (or click **Browse** to search for the file and location).

5 In the **Image Content** area:

➤ Select **Visible Window Only** to save only the part of the graph that appears in the window.

➤ Select **Draw Grid** to draw a grid in the graph

➤ Select **Selected CIs Only** to save only the CIs that you selected and their relationships.

6 In the **Image Characteristics** area:

➤ Enter a value in the **Image Quality (0-100)** window or move the indicator left or right to set the required quality of the printed map on the scale. 100 means that the quality is excellent.

7 In the **Size** area:
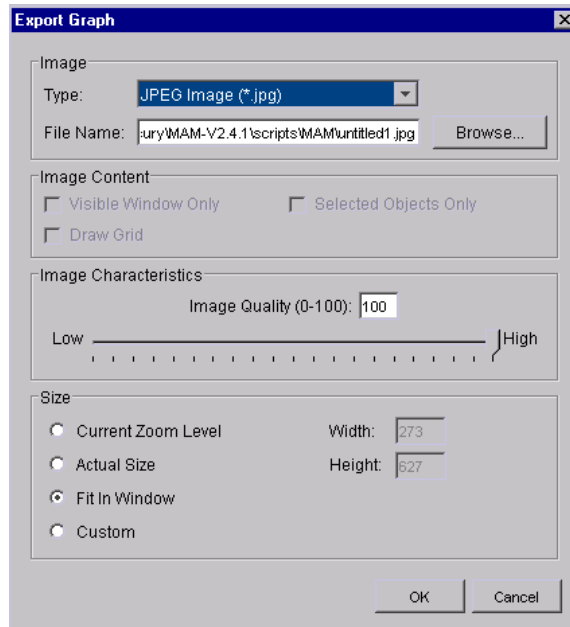
➤ Select **Actual Size** to save the view as it really is

➤ Select **Fit in Window** to expand or shrink the view so it fits in the window.

➤ Select **Custom** to specify the size of the view you want to save in the **Width** and **Height** boxes.

8 Click **OK**.

Printing the Topology Map

You can print the contents of the topology map, as well as the contents of the tabs in the Information pane. The result is similar to a screen capture. Therefore, it is recommended to arrange the contents of the topology map and Information pane tab according to your requirements, before printing.

This section includes the following topics:

➤ "Printing the Contents of the Topology Map in the Topology Map" on page 220

➤ "Printing the Contents of the Information Pane" on page 221

Printing the Contents of the Topology Map in the Topology Map

You can print the contents of the topology map in the Topology Map.

To print the contents of the Topology Map in the Topology map:

1 Select and display the contents you want to print in the Topology Map.

2 To arrange the contents of the topology map for print, you can use the **Layout**, **Layer Setup**, and **Hide** options.

For details about the **Layout** option, see "View Level" on page 40.

For details about the **Layer Setup** option, see "Improving Views" on page 197.

For details about the **Hide** option, see "View Shortcut Menu" on page 168.

3 Select **Map** > **Print Map** to open the Print Setup dialog box.

4 Select **Print Entire Graph** to print the complete view.

5 Select **Print Current Window** to print the part of the view that appears in the window.

6 Select **Print Current Selection** to print the selected notes and their relationships.

7 In the **Scale By** area:

➤ Select **Pages** and enter the number of columns and rows in the **Page Columns** and **Page Rows** boxes

➤ Select **Actual Size** to print the view as it really is.

➤ Select **Zoom Level** to print the view at the zoom level that you are currently using.

8 In the **Caption** area:

➤ Select **Print Caption**, enter the caption in the text area and click **Font** to select the font you want to use for the caption.

➤ Select the position of the caption in the **Position** list.

9 In the **Multipage Printing** area:

➤ Select **Print Page Numbers** if you want to print the page numbers.

➤ Select **Print Crop Marks** if you want to print them.

10 In the **Other** area:

➤ Select **Print Border** if you want the printed view to have a border and click **Color** to select the color you want for the border.

➤ Select **Print Background** if you want to print the background behind the view.

➤ Select **Print Grid** if you want to print a grid behind the printed view.

11 Click **Page Setup** to set the printer settings.

12 Click **OK**. These definitions are applied to your printouts. The Print dialog box appears.

13 Click **OK**.

Printing the Contents of the Information Pane

You can print the contents of the Information Pane.

To print the contents of the Information pane:

1 (Optional) Use the **Hide Column** and **Customize** options to arrange the contents of the Information pane for print. For details, see below.

2 Select **Map** > **Print** or click the **Print** button on the toolbar to open the Print dialog box.

3 Click **OK**.

Sorting, Hiding, and Displaying Columns

You can sort the contents of a column, hide a column, or display a hidden column. This section includes the following topics:

➤ "Sorting the Contents of a Column" on page 222

➤ "Hiding a Column" on page 222

➤ "Customizing a Display" on page 223

Sorting the Contents of a Column

You can sort the contents of a column.

To sort the contents of a column:

1 Click a column header.

2 To change the sort order, click the column header again.

Note: Once a column has been sorted its header displays a pink triangle pointed upwards for ascending order or downwards for descending order.

Hiding a Column

You can hide a column or display a hidden column.

To hide a column:

Right-click the column you want to hide and select **Hide Column**.

To display a hidden column:

Right-click anywhere in the **Configuration Item Browser** tab and select **Show All Columns**.

Customizing a Display

You can hide columns or display hidden columns.

To customize a display:

1 Right-click the header of a column in the **Configuration Item Browser** tab and select **Customize**. The Columns page opens.

2 Select columns in the **Hidden Columns** box or in the **Visible Columns** box and then use the left and right arrows to move the columns from one box to the other, or use the double left and right arrows to move all the columns from one box to the other.

3 Click **OK** to save your change.

Searching for CIs

Mercury Application Mapping searching capabilities enables you to search for CIs in the topology map, either in specific views and layers or in the entire CMDB. These capabilities also provide different search criteria, through which you can search for CIs according to their CIT, label, and/or attributes.

To search for CIs:

1 (Optional) If you want to search in a specific layer, open it in the Topology Map pane.

2 From the **Edit** menu, select **Find** to open the Find dialog box.

3 From the **Configuration Item Type** list, select the CIT of the CI(s) you want to find.

4 (Optional) Select **Derived** to find and display both the selected CIT(s) and any CIs derived from the selected CIT (inheritance CITs).

5 (Optional) In the **Options** area, if you want to find a specific CI(s) whose label you know, enter the CI's label in the **Label** box. Once you have entered a label, you can use the following search options:

➤ **Case Sensitive –** Distinguishes between uppercase and lowercase characters. When you select **Case Sensitive**, Mercury Application Mapping finds only those instances in which the capitalization matches the text you typed in the **Label** box.

➤ **Find whole words only –** Searches for occurrences that are whole words and not part of a larger word.

6 (Optional) In the **Advanced** area, to search for a CI(s) according to its attribute(s), do the following:

➤ Click the **Add** button to open the Condition dialog box.

➤ Define the attribute conditions according to the following table:

• **Attribute Name** – Select the required attribute from the list.

• **Operation** – Select the required operator (such as **=**) from the list (for details, see "Operator List" on page 44)

• **Value** – Enter or select the value of the attribute. (The box changes according to the attribute type you selected.) If the **Operator = In** or **Not In**, the **Value** string must appear in single quotes. If the **Operator = Like** or **Not Like** (for string only), the string **Value** must be surrounded by the percentage symbol.

• **AND/OR** – Enter **And** or **Or** to link multiple conditions.

Note: Attribute conditions cannot be defined for functional relationships.

➤ After you complete the attribute definition, you can change it by using the **Edit** and **Delete** buttons.

7 In the **Search in** area, select the place in which you want to conduct the search:

> ➤ **View –** Select from the **View** list the view in which you want to search.

> ➤ **My Views –** Searches only in the views that belong to your user profile.

> ➤ **Database –** Searches in the entire database.

8 (Optional) By default, **Get Full Path** is selected, and the view layers that appear above the CI(s) are displayed. Clear **Get Full Path** if you want to display only the view name in which the CI(s) is located.

9 Click **OK**. The search results are displayed in the Find Results tab in the Information pane.

10

Defining Instance Views

This chapter explains how to create an instance view. Instance views are not the result of TQL queries, but are views to which you manually attach CIs and relationships from the Mercury Universal CMDB or from an existing view.

This chapter describes:	On page:
What is an Instance View?	227
Creating an Instance View	227

What is an Instance View?

Instance views are views that are built by manually attaching individual CIs and relationships directly to the view. Instance views are not based on TQL queries. You can attach any CI from an existing view or from the CMDB into an instance view. If you choose to attach a CI from the CMDB, all existing relationships between the selected CIs are added to the view as well.

Creating an Instance View

You create a new instance view by selecting CIs from existing views or from the CMDB and attaching them to the new view. You can define additional view properties and also define organization rules to create a multi-level map. You can also perform a search for a CI or choose from a list of related CIs.

You can add a new instance view to any view folder. If you do not select a folder in the View Explorer, the new view is added to the Root folder.

This section includes the following topics:

➤ "Creating a New Instance View" on page 228

➤ "Displaying a CI's Attributes" on page 230

➤ "Searching for a CI" on page 230

➤ "Displaying Related CIs" on page 231

➤ "Deleting an Instance View" on page 232

➤ "Editing an Instance View" on page 232

Creating a New Instance View

This section describes how to create a new instance view and then attach the required CIs to the view.

Note: You do not need to attach any CIs to the view when you first create the view, but you must still click **OK** in the Instance View Definition window to create the view.

To create a new instance view:

1 In the View Explorer pane of the Topology View manager, select the folder in which you want to create the instance view.

2 Click **New** on the toolbar or open the **Map** menu and click **New** to open the Instance View Definition dialog box.

3 In the **View Name** box, type a unique name for the view.

4 (Optional) In the **Description** box, enter a description of the view.

5 (Optional) Click the **Properties** button to open the **View Properties** tab in the View Properties dialog box. Define additional view properties as described in "Creating a View With the Service View Manager" on page 110.

6 (Optional) Click the **Folding Rules** button to open the Relationship Rules dialog box. Define relationship rules to define multiple relationship rules. For details, see "Defining Multiple Relationship Rules" on page 151.

7 (Optional) To exclude specific links you do not want to appear in the view, do the following:

➤ Click the **Filter Links** button to open the Filter Links dialog box.

➤ Click the check boxes of the links you want to exclude from the view.

➤ Click **OK** to save the changes you have made.

8 In the left box, select the CI(s) that you want to attach to the view. You can select the CIs you want to attach to the instance view in the following ways:

➤ In the **Topology View** tab, select the required CI(s) at different levels of the view hierarchy. You can select multiple CIs by holding down the CTRL button and clicking the required CIs. To view and edit a CI's attributes, see "Displaying a CI's Attributes" on page 230.

➤ In the **Search** tab, search for CIs in existing views or in the entire CMDB. For details, see "Searching for a CI" on page 230.

➤ In the **Related CIs** tab, select a CI that is related to the selected CI in the **Topology View** tab. For details, see "Displaying Related CIs" on page 231.

9 Click the **Attach** button to attach the selected CIs to the view. The attached CIs appear in the **Configuration Items** box.

10 In the **Configuration Items** box, select the **Add SubGraph** check box if you want to attach all the descendents of the CI that exist in the CMDB (all the way down in the tree) to the view as well.

11 Click **OK** to create your view.

Note: It can take a few moments until the new view appears in the View Explorer pane.

Instance view and views created by TQL queries are represented by different icons, as described in the following table.

Icon	What it represents
🗁	A view based on a TQL query.
🖫	An instance view

Displaying a CI's Attributes

This section describes how to display and edit a CI's attributes.

To display and edit a CI's attribute:

1 In the Instance View Definition dialog box, select the **Topology View** tab.

2 Select the CI whose attributes you want to view.

3 Click the **Show CI's attributes** button to open the CI Attributes dialog box. The CI's Attributes dialog box enables you to view and edit the attributes defined for the selected CI. For details, see "Viewing and Editing CI Attributes" on page 201.

Searching for a CI

Mercury Application Mapping searching capabilities enables you to search for CIs to attach to your instance view either in specific views or in the entire CMDB. Theses capabilities also provide different search criteria, through which you can search for CIs according to their CI Type, label, and/or attributes.

To search for CIs:

1 In the Instance View Definition dialog box, select the **Search** tab.

2 From the **CI Type** list, select the CI Type of the CI(s) you want to find.

3 (Optional) Select **Derived** to find both CIs of the selected CIT in addition to any CI descending from the selected CIT.

4 (Optional) In the **Options** area, if you want to find a specific CI(s) whose
label you know, enter the CI's label in the **Label** box. Once you have entered
a label, you can use the following search options:

> ➤ **Case Sensitive –** Distinguishes between uppercase and lowercase
> characters. When you select **Case Sensitive**, Mercury Application
> Mapping finds only those instances in which the capitalization matches
> the text you typed in the **Label** box.

> ➤ **Find whole words only –** Searches for occurrences that are whole words
> and not part of a larger word.

> ➤ **Starts With –** Searches for occurrences that begin with the entered text.

> ➤ **Ends With –** Searches for occurrences that end with the entered text.

Note: The options **Starts With** and **Ends With** are only available if you have
not selected **Find whole words only**.

---

5 In the **Search In** area, select the place in which you want to conduct the
search:

> ➤ **My Views –** Searches only in the views that belong to your user profile.

> ➤ **Database –** Searches in the entire CMDB.

6 Click **Search**. The search results are displayed in the **Search** tab.

7 To do another search, click **Back to Search**.

Displaying Related CIs

Mercury Application Mapping enables you to display the interdependencies
of any selected CI (for details, see "Displaying Interdependent CIs" on
page 205). When you create a new instance view, you can choose to
populate the view with CIs from a list of related CIs.

You can display:

➤ All CIs connected to a selected CI that are part of a specific view.

➤ All CIs connected to the selected CI in the entire CMDB.

You select a CI in the **Topology View** tab, and the CIs related to the selected CI appears in the **Related CIs** tab.

To display the related CIs of a selected CI:

1 In the Instance View Definition dialog box, select the **Map View** tab.

2 Select the CI whose related CIs you want to display.

3 To display related CIs, do the following:

➤ Click the **Get Related CI from View** button to display related CIs in the selected view.

➤ Click the **Get Related CI from Database** button to display related CIs in the CMDB.

The related CIs are displayed in the **Related CIs** tab.

Deleting an Instance View

In the View Explorer, right-click the instance view you want to delete and select **Delete Instance View**.

Editing an Instance View

1 In the View Explorer, right-click the instance view you want to edit.

2 Select **Properties** to open the Instance View Definition dialog box.

3 Edit the instance view as required. For details, see "Creating a New Instance View" on page 228.

11

Event Management

This chapter introduces the Event system, which monitors the changes occurring in the managed world and delivers information about them.

Understanding the Event System

Mercury Application Mapping uses an event system to deliver information about the state of the managed world. This information begins as raw events, which are messages that notify Mercury Application Mapping about specific changes that have occurred in managed CIs. Three of the Mercury Application Mapping sub-systems—Discovery, Correlation, and Time Rules—determine which kind of changes that occur in CI attributes are registered as raw events.

Infrastructure		Application Server		Mercury Application Mapping GUI		Topology Map
Discovery Correlation Time Rules	Raw Events	CMDB	Event Messages (JMS events)	Event Configuration System (Rules Actions)	Active Events	Actions
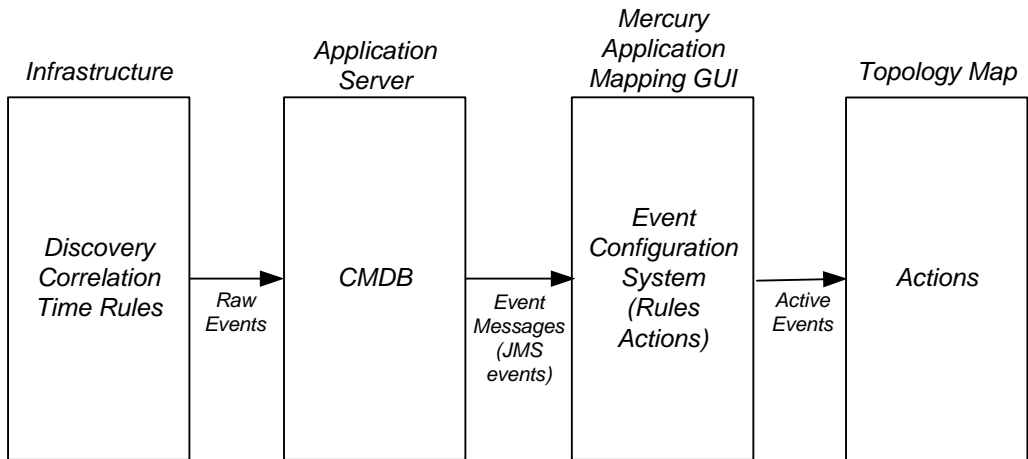
Raw Events

Raw events are saved to the CMDB. If a raw event is related to a CI that is displayed in the topology map, the raw event appears in the CI's Event log. For further details on the Event Log, see "Event Log" on page 249.

Severity	Correlation	Event Type	Message ...	Class Name	Create Time	Text Messa...	Discovery ...	Update Time
Normal	⊛	correlation...		nt	02/22/2004...	10.0.64.52-...	02/22/2004...	
Normal	⊛	correlation...		nt	02/22/2004...	10.0.64.52-...	02/22/2004...	
Minor(7)	⊛	correlation...		nt	02/22/2004...	ikbn,	02/22/2004...	
Minor(7)	⊛	correlation...		nt	02/22/2004...	ikbn,	02/22/2004...	
Critical	⊛	correlation...		nt	02/22/2004...	The lp cont...	02/22/2004...	
Normal	⊛	correlation...		nt	02/22/2004...	global_lab...	02/22/2004...	
Minor(7)	⊛	correlation...		nt	02/22/2004...	ikbn,	02/22/2004...	
Critical	⊛	correlation...		nt	02/22/2004...	The lp cont...	02/22/2004...	

Unfiltered — Ack: 0 — Total: 8

From the CMDB, the raw events are delivered, through JMS event messages, to the Event Configuration system. There, predefined event rules are applied to the raw events and determine whether these raw events affect the system and with what affect. Each event rule consists of at least one condition and at least one action. When a raw event meets one or more conditions, it triggers an action that is performed by Mercury Application Mapping.

Mercury Application Mapping enables administrators to define the rules that govern the way changes in the state of managed CIs are handled by the system. By defining these rules, administrators can determine which actions the system should take when changes that meet predefined conditions occur.

Note: For details on the configuration of the event system, see "Event Configuration" in the *Mercury Application Mapping Administration Guide*.

Example

The discovery system has determined that the administrative state of a certain CI has changed to **Unmanaged**. Consequently, it creates a new raw event that contains this information. The new raw event is sent as an event message to the Event Configuration System.

The Event Configuration System checks whether there are rules that can be applied to the new raw event. The following Event Configuration dialog box displays such a rule. The condition in the **Event Rules** list specifies that if an event message contains information about the administrative state of a CI, it should trigger an action that changes the CI's **admin** state in the Topology View.

Since the new raw event meets this rule's condition, the rule's action is performed, and the **admin** state of the CI is changed accordingly.
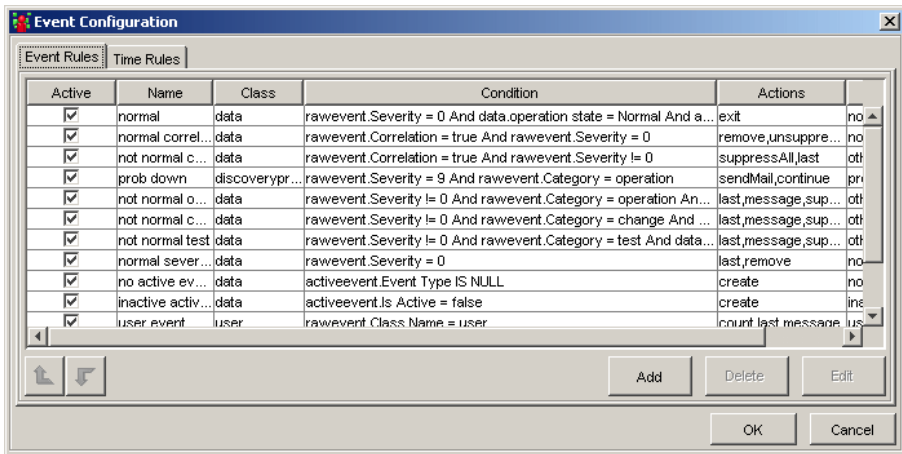
Although the Discovery, Correlation, and Time Rules systems continuously create raw events in response to changes occurring in the managed world, only those that meet at least one of the rule conditions are actually reported in Mercury Application Mapping. However, not all the raw events that meet the defined conditions are displayed in the Event browsers and manifested via state changes and icon blinks in the Topology View. (For details about Event Browsers, see "Displaying Event Browsers and Logs in Separate Windows" on page 248.) Only the raw events that are manifested and acknowledged are the ones that become active events.

The actions that are performed as a result of a condition fulfillment determine which raw events become active events and which remain as raw events. One action, which is specifically aimed at creating new active events from raw events, is called **Create.** Other actions that convert raw events into active events update the attributes of existing active events. By performing this update, they create new active events.

You use the Topology View to manage active events related to the CIs displayed in a selected view. This gives you the ability to locate problems that may be occurring in the system at the CI level. When a CI contains an event, except for an event that is a result of a correlation rule, the icon for the selected CI blinks and an exclamation mark is added before the icon in the View Explorer. You can then select the Configuration Item Events tab, or the other Event tabs and browsers, to view the event information and acknowledge the event:

237
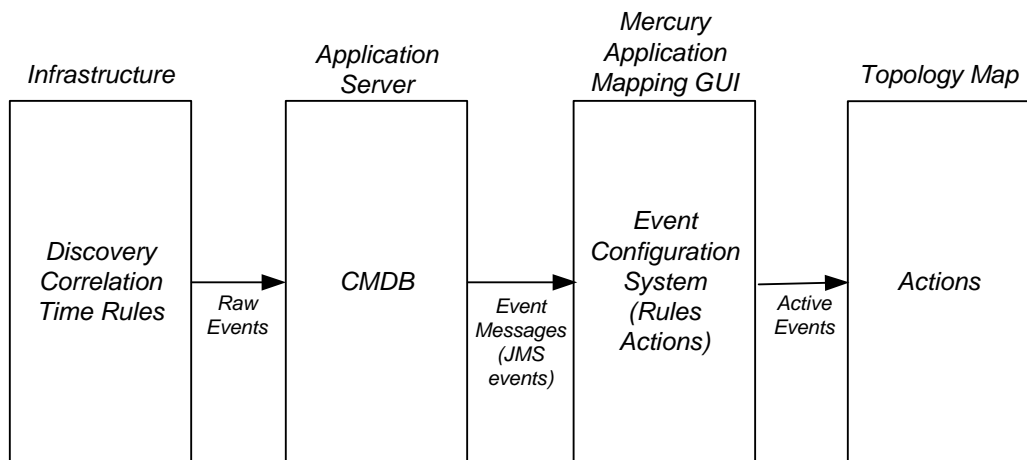
Once an event is acknowledged, the related icon on the Topology map stops blinking and the exclamation point disappears from the View Explorer.

Acknowledged event

Active events in the View Event button

Unacknowledged event
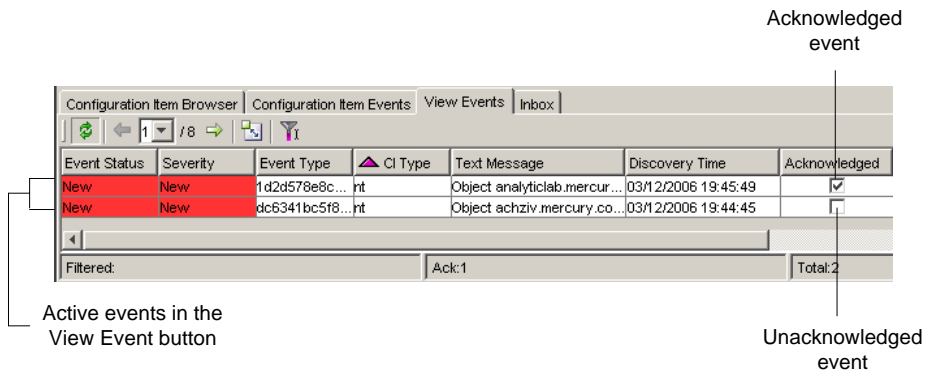
The Event Tabs in the Information Pane

This section describes the tabs in the Information pane, that enable you to manage events either at the CI or view level. For details of the attributes, see "Event Tab Definitions" on page 242.

Note: Event logs and browsers (regular and compound) display not more than 500 events.

This section includes the following topics:

➤ "Configuration Item Events" on page 239

➤ "View Events" on page 239

➤ "Inbox Tab" on page 239

➤ "Filtering the Events You Want to Display" on page 239

➤ "Determining the Number of Events That Appear on a Page" on page 241

➤ "Refreshing the Event Table" on page 241

➤ "Event Tab Definitions" on page 242

Configuration Item Events

The Configuration Item Events tab displays all the events related to the CI currently selected in the View Explorer or topology map. It includes information about the attributes of each listed event.

Click the **Compound** icon to toggle between showing and hiding the events of the selected CI's children.

View Events

The View Events tab displays all events related to CIs of a selected view in the View Explorer. It includes information about the attributes of each listed event.

Inbox Tab

The **Inbox** tab displays event messages that are sent to particular users or group members. Each message includes information about the event attributes.

Whenever the Inbox tab contains messages, an envelope icon is displayed in the status bar of the main window.

Filtering the Events You Want to Display

You use the Filter Settings dialog box to include only events you want to display in the Event tabs for the selected CI.

To filter the events you want to display:

 1 In the View Explorer or topology map, select a view or CI to which an event has been sent.

 2 Click the **Filter Settings** icon to open the Filter Settings dialog box.

 3 Select **is filter active** to display only the filter conditions defined in the Filter Settings dialog box. Clear the **is filter active** check box to display all events.

 4 To display events that were sent during a specified period of time:

 ➤ Select **Event creation time filter**.

 ➤ Click the date to the right of the **From** and **To** boxes to define the beginning and end of the time period.

➤ In the **Time** box, enter the required time.

➤ Click a date on the calendar. The times and dates you select are displayed in the **From** and **To** boxes.

➤ Click **OK**.

5 In the **Categories filter** area, select the required severity levels for the different categories.

6 From the **is Ack** list, select an option:

➤ **True** – display CIs whose events have been acknowledged

➤ **False** – display CIs whose events have not been acknowledged

➤ **None** – display CIs whose events are either in an acknowledged or unacknowledged state

7 From the **is Corr** list, select an option:

➤ **True** – display CIs whose events are the result of a correlation rule

➤ **False** – display CIs whose events are not the result of a correlation rule

➤ **None** – display CIs whose events are either the result of a correlation rule or not the result of a correlation rule

---

Note: Use any necessary combination of parameters you need to correctly define the filter conditions.

---

8 Click **Clear** to clear all the selected options.

9 Click **OK** to save the settings you have defined. The event list is filtered according to the filter settings you set.

Determining the Number of Events That Appear on a Page

The list of events are divided into pages. The number in the **Go to chunk** list indicates which page is currently being displayed. For example, 2/4 means that it is the second out of four pages.

To determine the number of events that appear on a page:

 1 Click the **Set bulk size** button to open the Set bulk size dialog box.

 2 Use the up and down arrows or type the number of events you want to appear on a page and click **OK**.

 3 To view other pages, use the left and right arrows.

Refreshing the Event Table

The **Refresh** button toggles between two states:

 ➤ The **Refresh** button, when in this state, automatically updates the event table with the latest information from the CMDB.

 ➤ The **Refresh** button, when in this state, indicates that the event table is not being updated automatically.

Event Tab Definitions

The following table describes the attributes of Mercury Application Mapping active events that make up the Configuration Item Events and View Events tab columns:

Note: Since an active event is always created from a raw event, it takes the same attributes as its raw event. In addition, active events have other attributes that are specific to them.

Column	Description
Severity	The severity of the event (Normal, Warning, Minor, Major, Critical). Each severity level is displayed in a predefined color.
Correlation	Indicates whether the event is the result of a correlation rule. When the event is a result of a correlation rule, a correlation icon appears in the box.
CI Type	The CIT of the CI for which the event was sent.
Event Type	The type of the event.
Message Group	Not in use.
CIT Name	The CIT name of the event's CI.
Create Time	Not in use.
Text Message	A textual description of the event.
Discovery Time	The date and time in which, a change in the state of a managed CI is discovered. This discovery caused the creation of the raw event that is related to the active event.
Update Time	Not in use.
Attribute Name	For events from Collector Threshold type only. Indicates the name of the attribute that caused the event.

Column	Description
Attribute Value	For events from Collector Threshold type only. Indicates the attribute value that caused the event.
Data 1-10	Undefined attributes that enable you to add your own event attributes.
System Time	The date and time of the entrance of the first related raw event to the database.
CI Label	The inclusive label of the CI related to the event. Includes other CIs that identify the selected CI. For example, a disk's label includes the disk name and its host name.
Rule Tracing	The names of the rules that cause the creation of the event and/or are involved in the event.
Acknowledged	Indicates whether the event has been acknowledged. In addition to changing the acknowledged state from the shortcut menu (for details, see "CI Shortcut Menu" on page 170), you can change it from this box. To acknowledge the event, select **Acknowledged**. To reverse the acknowledgment, clear **Acknowledged**.
Acknowledged By	The user by whom the event is acknowledged or unacknowledged after prior acknowledgment. This box is automatically selected once the event is acknowledged or unacknowledged.
Acknowledgment Time	The time when the event was last acknowledged or not.
Note	User notes. This box can be filled in by right-clicking the event and opening the Event Attribute dialog box. In the **Value** cell of the **Note** attribute you can enter your notes.
Last System Time	The last date and time the related raw event reached the database.
Last User Time	The last date and time in which, a change that cause the creation of the related raw event is discovered.

Column	Description
Counter	The number of times the same active event has arrived without any severity or message change.
Create Counter	The number of times the event have switched from an inactive state to an active one. (The switch from Active to Inactive is done by the **Clear** action, which deactivates the active event. As a result, the active event is removed from the Event Browsers. However, the event can still be restored, by each of the actions that update one of the active event's attributes.)
Replace Counter	The number of times the same active event has arrived with a severity change.
Active Time	The last date and time the event has become active.
Suppressed	For correlation events mainly. Determines whether the CI's icon blinks once the event occurs, similarly to the effect of the acknowledgment action. (Both actions determines whether the icon blinks. One is performed automatically by Mercury Application Mapping, while the other is performed by the user. If one of them does not apply to the CI, the icon still blinks.) The suppression action automatically defines, for all correlation rules, that CIs that are affected by a root-cause CI do not blink even though they have active events. For example, if a correlation rule is defined for the connection between a network CI and its hosts, the suppression action prevent the network's hosts from blinking once the network is down.
Event Status	The severity level of the CI, as adjusted for the weight according to its significance in the system. (The significance weight is set in the Status Factor tab in the View Node Definition dialog box. For details, see "Specifying the CI Status Factor" on page 141. For example, if the severity of the event is **5** and the status weight is defined at **100%**, the CI's Event Status is **5** (5*100/100). However, if the weight is defined as **60%**, the CI's status is **3** (5*60/100).

Column	Description
Category	The category for which the event is sent. For details on categories, see "Defining Attribute States" in *Mercury Application Mapping Administration Guide*.
Origin	Indicates the origin of the event. For example, a correlation rule, a user or a pattern.
Type	Gives additional information about the event. For example, the name of the time rule that triggered the event.

Event Shortcut Menu

The following table contains a brief description of each option in the shortcut menu, displayed by right-clicking an event in an **Event** tab or browser:

Option	Description
Delete	Deletes the selected event.
Note	Displays the Note dialog box. Enables you to add a note to a CI.
Show Event Attributes	Displays the attributes of the selected event. You can edit the values of the attributes that are displayed in italics.
Event Log	Displays a history list of all the raw events that occurred in relation to a selected CI.
Show Root-Cause	(This option is enabled only for events that are the result of a correlation rule.) Displays the Root-Cause Correlation Rules dialog box, enabling you to select a correlation rule, and then display in a new window the root-cause CI(s) that affected the selected relationship. If only one correlation rule is defined for this CI, the Correlation window is directly displayed.
Show Configuration Item Attributes	Displays the attributes of the CI related to the selected event. You can edit the values of the attributes that are displayed in italics.
Event Navigator	Selects the node to which the event was sent in the topology map and View Explorer.

Sorting and Customizing Event Tabs

In each tab, you can sort the displayed events and set which columns to show and in what order.

This section includes the following topics:

➤ "Sorting an Event Tab" on page 247
➤ "Hiding a Column in an Event Tab" on page 247
➤ "Displaying Hidden Columns in an Event Tab" on page 247
➤ "Customizing an Event Tab" on page 248

Sorting an Event Tab

You can sort an Event tab.

To sort an Event tab:

On the selected tab, click the column header tab. To change the sort order, click the column header again.

Hiding a Column in an Event Tab

You can hide a column in an Event tab.

To hide a column in an Event tab:

On the selected tab, right-click the header of the column you want to hide and select **Hide Column**.

Displaying Hidden Columns in an Event Tab

You can display hidden columns in an Event tab.

To display hidden columns:

On the selected tab, right-click one of the column headers and select **Show All Columns**.

Customizing an Event Tab

You can customize an Event tab.

To customize an Event tab:

1 On the selected tab, right-click one of the column headers and select **Customize** to open the Columns dialog box.

2 To remove column(s) from the tab, select the required column(s) in the **Visible Columns** area and click the **Remove Column** button. The selected column(s) moves to the **Hidden Columns** area. (To return a column to the **Visible Columns** area, select it and click the **Add Column** button.)

3 To change the display order of the columns, use the up and down arrows.

4 Click **OK** to apply your customization to the tab.

Displaying Event Browsers and Logs in Separate Windows

In addition to viewing event information in the Event tabs, you can display event browsers and log in separate windows:

➤ Event Browser. For details, see "Event Browser" on page 248.

➤ Event Log. For details, see "Event Log" on page 249.

➤ Compound Event Browser. For details, see "Compound Event Browser" on page 249.

➤ Compound Event Log. For details, see "Compound Event Log" on page 250.

Event Browser

The Event browser displays the same information contained in the Configuration Item Events tab in a separate window.

To display the Event Browser in a separate window:

Right-click a CI in the View Explorer or Topology View and select **Events** > **Event Browser** to open the Event Browser window.

For details about the information displayed, see "Event Tab Definitions" on page 242.

Event Log

The Event log displays a list of raw events that occurred to a selected CI. Through this log, you can view the event history of the CI.

To display the Event log:

Right-click a CI in the Explorer or Topology map and select **Events** > **Event Log** to open the Event log.

All the columns that appear in the Event Log, except the **parameter** column, are the same as those that appear in the Event tabs and browsers. For details about the information displayed, see "Event Tab Definitions" on page 242.

The **parameter** column can contain two types of data. The first data type is displayed when the raw event is related to an action that is performed by a user. The data consists of a textual description of the action (for example, **Send Event**, **Ack**) and the name of the user that performed it. The second data type is displayed when dynamic parameters are defined in the system's Discovery Patterns. These dynamic parameters can be sent as part of the raw event values, and they can be addressed and used by the Event Rules.

Compound Event Browser

The Compound Event browser displays the information contained in the Event Browser tab, as well as event information about the selected CI's children (if children exist).

To display the Compound Event browser:

Right-click a CI in the View Explorer or topology map and select **Events** > **Compound Event Browser** to open the Compound Event Browser.

For details about the information displayed, see "Event Tab Definitions" on page 242.

Compound Event Log

The Compound Event log displays the information contained in the Event Log tab as well as event log information about the selected CI's children (if children exist).

To display the Compound Event log:

Right-click a CI in the Explorer or Topology map and select **Events** > **Compound Event Log** to open the Compound Log browser.

Receiving Change Events

This section describes the Change event type. It includes:

➤ "Receiving Change Events When CIs Are Added or Deleted" on page 250

➤ "Receiving Change Events When a Change Monitored Attribute is Modified" on page 251

Receiving Change Events When CIs Are Added or Deleted

Mercury Application Mapping enables you to receive events of the type Change when:

➤ CIs are added to maps

➤ CIs are deleted from maps

To receive events of type Change when CIs are added to, or deleted from maps:

1 In the Service View Manager, right-click the required CI and then click **Properties** to open the **Properties** tab in the View Properties dialog box.

2 In the **Send notification on changes** section, select:

➤ **Added CIs** - if you want to be notified when CIs are added

➤ **Removed CIs** - if you want to be notified when CIs are removed

3 Click **OK** to save the changes you have made. An event is sent every time a CI is either added to or removed from maps.

Receiving Change Events When a Change Monitored Attribute is Modified

Mercury Application Mapping enables you to track modifications made in any CIT attribute defined as **Change Monitored**. The changes that occur appear in the Information pane.

To receive change events when a Change Monitored attribute is modified:

1 In the CI Type Manager, right-click the required CI and then select **Edit CIT**.

2 Click the **Attributes** tab and then click **Add** to open the Add Attribute dialog box.

3 In the **Configuration management only** area, ensure that **Change Monitored** is selected for the selected attribute.

4 Click **OK**.

5 In the Service View Manager, right-click the required CI and click **Properties** to open the View Properties dialog box.

6 Click the **View States** tab

7 To view change events and statuses on the maps themselves, add the Change state. For details, see "About Defining Topology Map States" on page 113.

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

Every time a change occurs, a change event is sent to the view. If the Change Monitored attribute of the CI has changed, the name of the CI that has been changed is displayed.

Sending an Event

This section describes how to define the event category that affects a CI.

To send an event:

1 Right-click a CI in the Topology View and select **Actions** > **Send Event** to open the Send Event dialog box.

2 From the **Category** list, select the category that is affected by the event. The severity list that is displayed is created in the System Type Manager. For details, see "Enumerations and Lists" in the *Mercury Application Mapping Administration Guide*.

3 Select the required severity.

4 Click **OK**. The color and blinking of the CI are changed when the event is sent under the conditions defined in the **State Priority** list. For details, see "About Defining Topology Map States" on page 113.

Clearing All Events

This option enables you to delete the events belonging to a CI of a certain category or of all categories.

To clear the events from a CI:

1 In the Topology Map, right-click the CI from which you want to clear events to open the Clear Events dialog box.

2 Choose an option as follows:

➤ **Category** – Delete events belonging to a specific category. Select the required category. The list contains the categories that you defined in the Enumeration Manager.

➤ **All Categories** – Delete the events belonging to all categories.

> **Note:** When clearing all events for a specific category, the CI may still be affected by other category events. Only if you choose to delete the events for all category types, does the CI's state return to Normal.

3 Click **OK** to save the settings you have defined.

Part V

Generating Configuration Management and Change Reports

12

Configuration Management Reports

This chapter describes the reports you can generate in Mercury Application Mapping.

Generating an Asset Report

Mercury Application Mapping enables you to create a report that lists all the CIs in a selected view and their attribute values.

To generate an Asset Report:
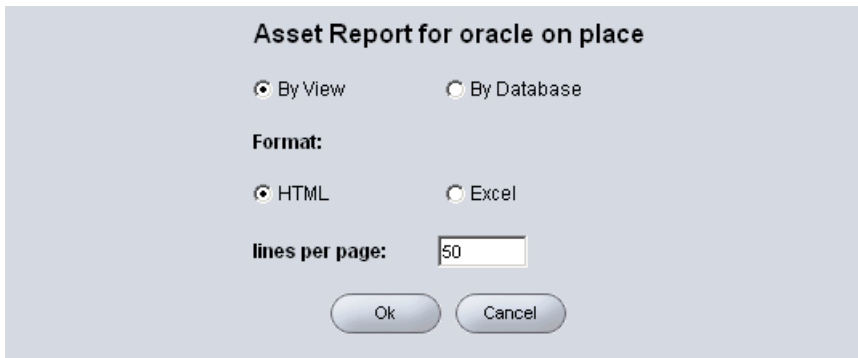
1 In the Topology Map, right-click the view for which you want to create an Asset Report.

Note: The report contains only the attributes whose **Asset Data** check box was selected in the **Attributes** tab in the CI Type Manager. For more details, see "Creating CITs" on page 384.

---

2 Select **Reports > Asset Report** to open the following in your browser.

3 To define what you want the report to include:

> ➤ Select **By View** to include only CIs contained in this view.

> ➤ Select **By Database** to include additional information from the CMDB. By choosing this option, the report will include CIs that are connected by the **container_f** relationship to the CIs in selected view. For example, the report would include the CPU and memory of a server that is contained in the view.

4 Define the format in which you want to receive the report. You can receive the report in the following formats:

> ➤ **HTML**

> ➤ **Excel**

5 If you select **HTML**, enter a value in the **lines per page** box to define the maximum number of rows you want to appear on each page.

6 Click **OK** to generate the report.

The Asset Report retains the same hierarchical structure as in the Topology Map.

Generating a Dependency Report

This report is based on TQLs that are created in the Enrichment Manager that use the relationship of the type **dependency**. Mercury Application Mapping enables you to create a report that lists all the relationships of the type **dependency** that connect between servers in the selected view or between the servers in the CMDB and the servers within the selected view.

The report includes relationships that:

➤ Exist between databases and clients

➤ Exist between clients and servers

➤ Are created as a result of an MQ connection

➤ Are created as a result of a P2P connection

To generate a Dependency Report:

1 In the Topology Map, right-click the view for which you want to create a Dependency Report.

2 Select **Reports** > **Dependency Report** to open the following in your browser.

3 To define what you want the report to include:

➤ Select **By View** to display all the dependencies between the servers in the selected view.

➤ Select **By Database** to display all the servers in the Mercury Universal CMDB that are dependent on the servers in the selected view.

4 Define the format in which you want to receive the report. You can receive the report in the following formats:

➤ **HTML**

➤ **Excel**

5 Click **OK** to generate the report.

Host	Link	Counter	Clients
DANGEROUS	Oracle:skazal	1	DANGEROUS
ZEPPO	nbsession:139	1	MILLER
ILLABBDC01	nbsession:139	1	MILLER
SAMSON	nbsession:139	2	INFRADC,MILLER
MILLER	netbiosW2K:445	1	place.mercury.co.il
	nbsession:139	1	FRIENDS
	sql:1433	4	BUDWISER,HAMPSHIRE,LEFF,MAMBO
INFRADC	netbiosW2K:445	2	place.mercury.co.il,KALI
EXPEDITION	nbsession:139	2	WHISTLE,DR2
PROWLER	Oracle:skazal9	1	PROWLER
DR2	netbiosW2K:445	2	place.mercury.co.il,MUST
CALYPSO	nbsession:139	1	MILLER
PLANET	sql:1433	1	INFRADC
MURPHY	nbsession:139	2	INFRADC,MILLER
WHISTLE	netbiosW2K:445	7	place.mercury.co.il,MARIACHI,TAS2,host,host,host,host
	http:80	4	JOOP,host,host,host
	oracle:1521	1	WHISTLE

The report contains the following columns:

Column	Description
Host	The servers in the selected view.
Relationship	The type of relationship that exists between the server and their clients.
Counter	The number of clients to which the server is connected.
Clients	The list of clients that are linked to the servers.

Note: You can create new TQLs in the Enrichment Manager using the **dependency** relationship to describe the connection between hosts. Every new **dependency** relationship that is used appears in the Dependency Report.

Generating an Event Report

Mercury Application Mapping enables you to generate a report displaying the events for the CIs in the selected view. You can define which events you want included in the Event Report.

To generate an Event Report:

1 Right-click the view for which you want to create an Event Report.

2 Select **Reports** > **Event Report** to open the following in your browser.

3 To display events that were sent during a specified period of time, do the following:

> ➤ Click the button to the right of the **From** and **To** fields to define the beginning and end of that time period respectively.

> ➤ In the **Time** box, enter the required time.

> ➤ Click a date on the calendar.

The times and dates you have selected appear in the **From** and **To** boxes.

4 Select the required severity levels for the different categories.

5 From the **isAck** list, select one of the following options:

> ➤ Select **True** to display CIs whose events have been acknowledged.

> ➤ Select **False** to display CIs whose events have not been acknowledged.

> ➤ Select **None** to display CIs whose events are either in an acknowledged or unacknowledged state.

6 From the **isCorr** list, select one of the following options:

> ➤ Select **True** to display CIs whose events are the result of a correlation rule.

> ➤ Select **False** to display CIs whose events are **not** the result of a correlation rule.

> ➤ Select **None** to display CIs whose events are either the result of a correlation rule or not the result of a correlation rule.

7 From the **Sort By** list, select a category by which you want to sort the event list.

8 Click **OK** to generate the Event Report.

The Event Report has the following fields:

Field	Description
Severity	The severity of the event. Each severity level is displayed in a different color.
Category	The category for which the event was sent.
Configuration Item Type	The CIT name of the event's CI.
Configuration Item	The CI's label.
Text Message	A textual description of the event.
Ack	Indicates whether the event has been acknowledged or not.
Discovery Time	The date and time in which the event was triggered.
Corr	Indicates that the event occurred as a result of a correlation rule.
Origin	Indicates the origin of the event. For example, a correlation rule, a user, or a pattern.

13

Change Report

This chapter describes how to generate a Change report in Mercury Application Mapping.

This chapter describes:	On page:
Understanding Change Reports	267
Generating a Change Report for a CI	268
Performing a Snapshot Comparison	271
Viewing the Change History for a CI	271

Understanding Change Reports

The Change report displays information about the changes made to the properties of CIs, which were assigned to keep this information. The report displays changes that occurred in a selected CI and all its children.

A CI is displayed in the Change report only when the following conditions occur:

➤ At least one of the history-sensitive properties of the CI has been updated.

➤ The property that was changed had been assigned the **Change Monitored** attribute in the last 24 hours. For details, see "Creating CITs" on page 384.

This section includes the following topics:

➤ "Generating a Change Report for a CI" on page 268

➤ "Performing a Snapshot Comparison" on page 271

➤ "Viewing the Change History for a CI" on page 271

Generating a Change Report for a CI

This section describes how to generate a Change report within a specific time frame. The Change report displays the changes that occurred within a defined time frame.

Note: The Change report option is also available for any CI or CI's children whose history-sensitive properties were modified.

This section has the following topics:

Generating a Change Report

You can generate a Change Report.

To generate the Change report for the selected CI its children:

1 In the Topology View, right-click a CI whose properties have changed.

2 Select **Reports** > **Change Report** to open the Change Report Time Frame.

3 To define the required time frame for the Change report, do the following:

➤ Click the buttons to the right of the **From** and **To** fields to define the beginning and end of the required time period.

➤ In the calendar that opens, choose the required date and time.

➤ Select **Now** to define the end of the period as the current date.

4 Click **Generate** to display the Change report for the selected CI and its children.

A typical Change report is as follows:

The following information is displayed in the report:

Item	Description
CI Label	The name of the CI.
CIT Name	The name of the CIT (internal).
CI Attribute	The name of the CI attribute that was updated.
New Value	The new value of the CI attribute.
Old Value	The previous value of the CI attribute (before the change).
Change Date & Time	The date and time when the update was performed.

Item	Description
Changer	The name of the user that manually modified the CI's property, or the name of the discovery probe that automatically discovered a change made to the CI's property.
History	Click the icon to view the change history for the CI. For details, see "Viewing the Change History for a CI" on page 271.
Removed/Added CIs	The number of CIs that were removed/added to the view.
Removed/Added Relationships	The number of relationships that were removed/added to the view.

Sorting the Contents of a Column in the Report

You can sort the contents of a column in the report.

To sort the contents of a column:

1 Click a column header.

2 To change the sort order, click the column header again.

Once a column has been sorted, its header displays a triangle pointed upwards for ascending order or downwards for descending order.

Viewing the Report in Excel Format

You can view the report in Excel format.

To view the report in Excel format:

Click the **Export to Excel** button at the top-left corner of the Change Report.

Printing the Contents of the Report

You can print the contents of the report.

To print the contents of the report:

Click the **Printer Friendly** button at the top-left corner of the Change Report.

Performing a Snapshot Comparison

Mercury Business Availability Center enables you to perform a snapshot comparison if snapshots for the selected view were saved in Mercury Application Mapping. For details about snapshots, see "About Environment Comparison" on page 406.

To perform a snapshot comparison:

 1 At the bottom of the Change Report window (for details, see "Generating a Change Report for a CI" on page 268), click **View Snapshots** to open a separate window in your Web browser displaying the snapshots that were saved for the selected view.

Note: The View Snapshots option appears only if snapshots for the selected view were saved.

---

 2 Select the two snapshots you want to compare.

 3 Click **Compare** to open a separate window in your Web browser. The window is divided into two sections. Each side displays a hierarchical tree structure of all the CIs in the topology map for that snapshot version.

Viewing the Change History for a CI

You can view the change history of a specific CI.

Note: To display the Change History report, make sure that popups are not blocked in your browser.

To view the change history for a CI:

1 Double-click the **History** icon for the required CI to open the History page.

A typical Change History report is as follows:

The page displays the following information:

Item	Description
CI Attribute	The name of the CI attribute that was updated.
New Value	The new value of the CI attribute.
Change Date & Time	The date and time when the update was performed.
Changer	The name of the user that manually modified the CI's property, or the name of the discovery probe that automatically discovered a change made to the CI's property.

➤ To sort the columns of the report, see "Sorting the Contents of a Column in the Report" on page 270.

➤ To view the report in Excel format, see "Viewing the Report in Excel Format" on page 270.

➤ To sort the columns of the report, see "Printing the Contents of the Report" on page 270.

Part VI

Correlation Rules

14

Using the Correlation Manager

This chapter introduces the Correlation Manager, which enables you to create topology correlations.

This chapter describes:	On page:
Quick Tour of the Correlation Manager	277
Correlation Manager Shortcut Menus	280
Working with the Correlation Manager	282

Quick Tour of the Correlation Manager

Mercury Application Mapping uses topology knowledge gathered through TQL queries to automatically outline the interdependencies that exist between service components in a specific IT infrastructure. By using this information on application and management dependencies, the root causes of system alerts can be discovered. In addition, the implications and consequences of any action, update, new deployment, or a failure of the service delivery and customer experience can be evaluated and handled accordingly.

By creating correlation rules based on TQL queries, you can determine the true origin of a problem, its business impact, and its appropriate prioritization. You can also establish automatic responses to recurring problems.

The Correlation Manager, which is displayed by selecting the Correlation Manager tab in the View Explorer, enables you to define correlation rules, which specify a casual relationship between two or more correlation nodes – nodes and physical relationships. Each correlation rule you define is attached to a specific Correlation TQL query. This type of TQL query is created especially for correlation rules, and it provides them with the required dependency information for their execution.

Once the casual relationship is established, the Correlation Manager enables you to determine the effect of the events occurring in root cause nodes. You can define the type and format of correlation events that are generated by root cause nodes, how they change the Correlation and Admin states of the affected nodes, and who is notified of the correlation occurrence.

A change in the correlation state of an affected CI can alter the CI's map status, which is manifested through a change in the color of the CI's symbol. This happens when the correlation event's severity is higher than other related events' severity. However, even if the symbol's color changes, the symbol does not blink, unlike all other cases in which the CI contains unacknowledged events. This is caused by a built-in utility, which automatically suppresses the blink of an affected CI's symbol when it is caused by the occurrence of a correlation event. It is designed to help you differentiate between affected and root cause CIs, while you determine and prioritize your responses to system alerts.

For example:
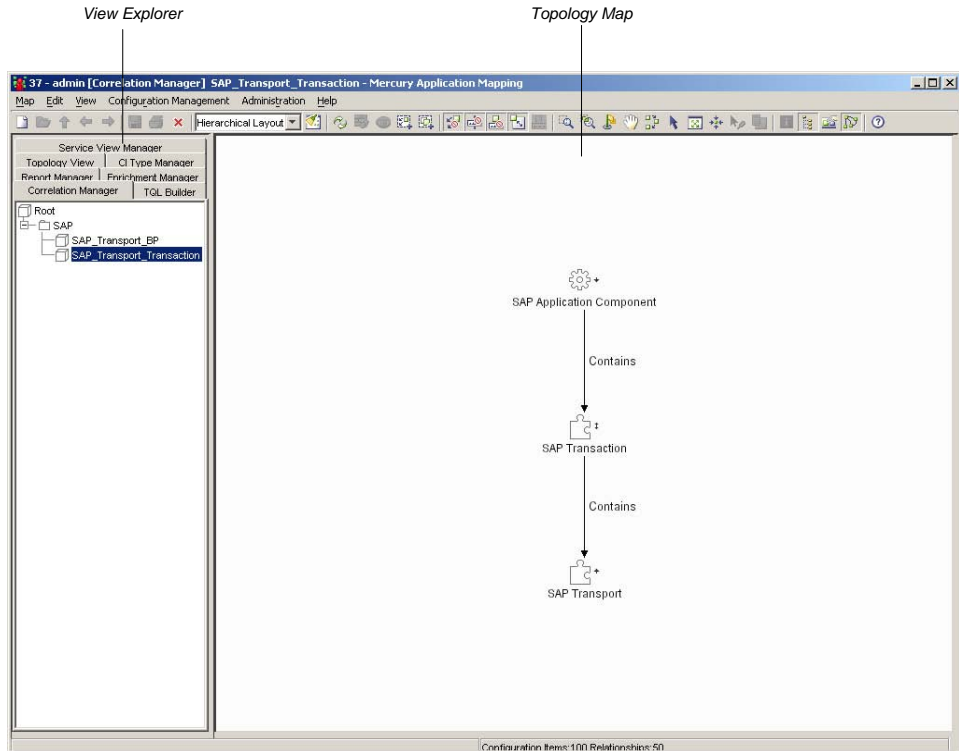
Correlation events are displayed in the **Configuration Item Events**, **View Events**, and **All Event** tabs in the Information pane of the Topology View. An icon appears in the **iscorr** column, when the event is a result of a correlation occurrence. The relationship between root cause CIs and affected CIs can be examined in the Topology Map via the **Show Impact** and **Show Root Cause** options, as described "Showing a Correlation Impact" on page 212 and "Showing Root Cause CIs" on page 214.

When the Correlation Manager tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **View Explorer –** Displays a hierarchical tree structure of the correlation rules you have defined. If required, you can export correlation rules and save them as XML scripts, which can be used for backup purposes.

➤ **Topology map –** Displays the currently selected rule, which consists of nodes that are defined in the TQL query and the relationships between them. The trigger node is marked by an up arrow that is located to the right of the node's icon. Affected nodes are marked by a down arrow to the right of the affected node's icon. Nodes that are both affected and trigger nodes are marked by a double-sided arrow.

Correlation Manager Shortcut Menus

The Correlation Manager contains different right-click shortcut menus, depending on your selection.

This section includes:

➤ "Correlation Rule Shortcut Menu" on page 280

➤ "Node/Relationship Shortcut Menu" on page 281

Correlation Rule Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a selected correlation rule in the View Explorer:

Option	Description
New	Enables you to define a new correlation rule.
New Folder	Enables you to create a new correlation folder.
Save	(Enabled only when a new Correlation query is created or when changes are made to an existing one.) Enables you to save the correlation rule to the CMDB.
Save As	Displays the Save As dialog box, enabling you to define a new name and description for a query. This option can be used for copying queries.

Option	Description
Delete	Enables you to delete the selected correlation rule from the database.
	Note that deleting a correlation rule deletes all events related to the rule.
Export	Displays a standard Save As dialog box, enabling you to save the correlation rule as an XML script. This option can be used to move correlation rule from one workstation to another, provided the related TQL query is also relocated.
Properties	Enables you to change the description and activation state of the rule.

Node/Relationship Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a correlation node or a relationship in the topology map:

Option	Description
Straight	(Enabled only for relationships with angles on the Topology map.) Enables you to straighten relationships with angles.
Define Affected	Displays the Root Cause dialog box, enabling you to define the affected node(s) and the correlation events that are generated by events in the root cause node.
Reset Affected	Removes all defined affected nodes.

Working with the Correlation Manager

Use the Correlation Manager to perform the following functions:

➤ Create new correlation rules, as well as duplicate and delete existing rules.

➤ Define which nodes are affected by the root cause node.

➤ Define the correlation events that are generated by events occurring in the root cause node.

➤ Define the format of messages sent to nodes and users regarding correlation events.

➤ Show or hide selected nodes and relationships in the topology map.

➤ Export and import selected correlation rules from XML scripts. This can be used to relocate correlation rules from one workstation to another providing the related TQL query is also relocated.

➤ Zoom in and out of the topology map area to view the selected layer at different levels of magnification.

➤ Select different layout views.

15

Defining Correlation Rules

This chapter explains how to define correlation rules.

About Defining Correlation Rules

When creating a correlation rule, you first name the rule and attach it to a specific Correlation TQL query. For details, see Part II, "Topology Query Language Builder." You then proceed to define which node in this scenario is the root-cause node, and what nodes are affected by it. This includes defining the actions that are performed on these affected nodes, as well as the messages that are sent to particular users (or groups of users) about correlation events.

Correlation Rule Workflow

You create correlation rules according to the following workflow:

➤ Create a correlation rule (for details, see "Creating a Correlation Rule" on page 284).

➤ Define which nodes and users are affected (for details, see "Defining Affected Nodes and Users" on page 285).

➤ Save the correlation rule (for details, see "Saving the Correlation Rule" on page 291).

Creating a Correlation Rule

The first step when creating a correlation rule is to define a name and description for the rule, attach the rule to a specific Correlation TQL query, and define whether the rule should be active in the system from the moment it is saved.

Note:

➤ Any major changes made to the TQL query after creating a correlation rule causes the system to delete the rule. These changes include deleting a root-cause or affected node or changing the Min and Max relationship definitions (for details, see "Defining Attribute Conditions for Nodes and Relationships" on page 126).

➤ Minor changes, such as adding a node to the TQL, do not cause the deletion of the rule.

➤ The correlation TQL that serve as the basis of the correlation rule, should comply with the restrictions described "Validation Restrictions" on page 93. If the TQL is not valid, it cannot be used for the creation of a correlation rule.

To create a correlation rule:

1 Click **New** button on the toolbar to open the New dialog box.

2 In the **Name** box, enter a unique name for the correlation rule.

3 (Optional) In the **Description** box, enter a description of the rule.

4 From the **Attached TQL** list, select the correlation TQL to which the correlation rule should be linked.

For details on creating TQL queries, see "Defining Topology Query Language (TQL) Queries" on page 57

5 (Optional) Verify that **Active** is selected to activate the rule in the system as soon as it is saved.

Note: If the rule is not activated now, it can be activated later by right-clicking it in the View Explorer, selecting **Properties** from the shortcut menu, and selecting **Active**.

---

6 Click **OK**. The new correlation rule is displayed in the tree in the View Explorer on the left side of the main window and in the Topology map on the right side of the window. For details of the Mercury Application Mapping main window, see "Mercury Application Mapping Main Window" on page 21.

Defining Affected Nodes and Users

The next step in the procedure is to define which node in the TQL query is the root-cause node, that is, the trigger of the correlation event, and which nodes are affected by this root-cause. (This information should be known beforehand, when first creating the TQL query, but the definition is performed here.)

This definition includes several stages:

➤ Defining correlation conditions for the attributes of the root-cause node

➤ Defining how many instances of the root-cause node should fulfill the conditions for activating the correlation actions

➤ Defining the nodes that are affected by the root-cause node and the users/groups that are to be notified once correlation conditions are fulfilled

➤ Defining the actions that are to be performed once correlation conditions are fulfilled

When selecting nodes to function as correlation triggers, they should comply with the following restrictions:

➤ You can select more than one node as a trigger. However, you cannot define a node as affected and as a trigger.

➤ If a node has a relationship whose minimum limit is **0** (meaning that one of its ends does not necessarily have a node linked to it), the node that is linked to its other end cannot be a root-cause node (since it may or may not exist in the TQL). For details about minimum limits, see "Defining TQL Nodes and Relationships" on page 124.

For example, object_15 cannot be either a root-cause or affected node because it is connected to the host with a Min limit of **0**.
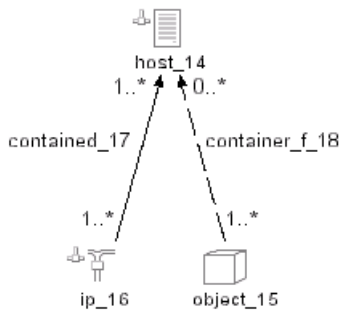
A node that is not visible cannot be a root-cause or an affected node.

To adjust the correlation rule definition to a management environment that contains multiple categories, Mercury Application Mapping enables you to specify the category that triggered the correlation rule and the event category that is sent to the affected nodes.

This section includes the following topics:

➤ "Setting Multiple Conditions for Defining Affected Nodes" on page 287

➤ "Editing a Correlation Rule" on page 291

➤ "Removing a Correlation Rule" on page 291

Setting Multiple Conditions for Defining Affected Nodes

You can set multiple conditions for defining affected nodes.

To set multiple conditions for defining affected nodes:

1 In the Correlation Manager, right-click the node or relationship that functions as a trigger.

2 Select **Define Affected** to display the affected nodes.
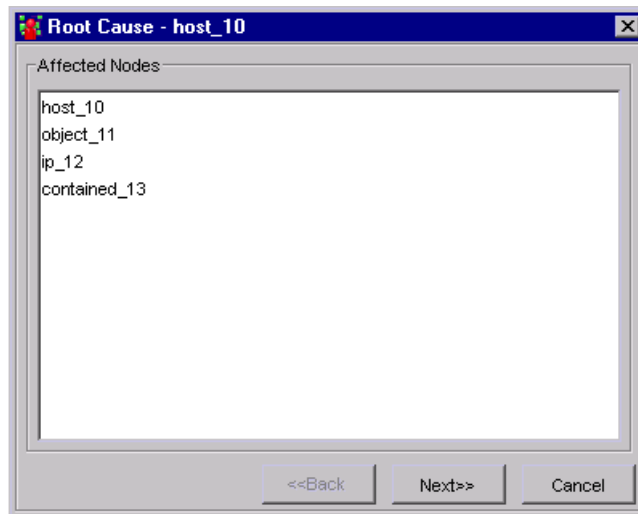
3 Select the node or nodes to be affected by the trigger.

4 Click **Next** to add a correlation rule.

5 Click **Add** to define a correlation rule.

6 In the **Description** box, type a description of the correlation rule you are defining.

7 In the **Conditions** area, click **Add** to define conditions for the node's attributes. When these conditions are met, the correlation actions are triggered.

For example, you can determine a correlation condition that states that the Operational State of the node is other from Normal. When the Operational State of the node instances changes and is no longer Normal, the condition is met and the assigned actions are activated.

Tip: It is recommended to define conditions for non-calculated attributes (for example, CIT or Names) in the TQL Builder and conditions for calculated attributes (such as States) in the Correlation Manager.

---

8 In the Condition dialog box that is displayed:

➤ Define the attribute conditions and click **OK**. For details on defining attribute conditions, see "Setting TQL Node and Relationship Definitions" on page 68.

Note: You cannot create a correlation rule using the **LIKE** or **NOT LIKE** operations.

➤ To define more than one condition, click the **Add** button again.

➤ Click **OK** to save your changes.

9 In the **Scope** area, define how many instances of the root-cause node should fulfill the conditions for activating the correlation actions.

For example, you can define that a correlation event is sent only when at least 10% of the system's hosts are down.

You can apply the trigger conditions not only to one of the instances of a root-cause node, but also to a certain percentage of the root-cause instances or to all of them. In addition, if you use the Send (Range) Correlation Event action, you can define whether numerous correlation events are sent, one for each root-cause instance (Any), or whether one event only is sent for the entire instance group (All, Range).

Select one of the following options:

➤ **Any –** One instance or more should fulfill the conditions. A correlation event is sent for each instance.

➤ **All –** All instances should fulfill the conditions. One correlation event is sent for all instances.

➤ **Range –** A certain percentage of all instances should fulfill the conditions. One correlation event is sent for all the instances (of the root-cause node) that are included in the range. For example, if you enter a range of 50% to 100%, the actions are triggered when 50% or more of the instances meet the conditions.

10 In the **Send Correlation Event** area:

➤ Select the category type for which you are sending the event in the **Category** list.

Important: The category you choose must be the same category you selected in the Condition statement. For example, if the Attribute Name you defined in the condition statement is **Performance State**, you must select the same category from the **Category** list, that is, **Performance**.

➤ In the **Message** box, define the message to be generated by the system, as follows:

• Click **Format** to open the Format Message dialog box.

• Select the attribute(s) to appear in the message by clicking the plus button next to each attribute. The selected attribute is displayed in the **Format** box.

- If required, you can use the function buttons to display multiple attributes in the label.

- Click **OK**. The defined message is displayed in the **Message** box of the Root-Cause dialog box.

➤ In the **Message** box, enter additional text to the message, as required.

11 In the **Event Severity** area, define the severity of the correlation event, by selecting one of the following options:

➤ Select **Fixed severity** to define a fixed severity level for every correlation event generated then select the required severity level in the list.

➤ Select **Severity relative to trigger severity (%)** (displayed when you selected the **Any** option) to define a severity level for each correlation event that is relative (as a percentage) to the severity of the trigger event. Enter the percentage in the box provided.

For example, if the severity of the trigger event is Critical(10) and this field is set to 80%, then the severity of the correlation event is Major(8).

➤ Select **Function** (displayed when you selected the **All** and **Range** options) to define a severity level for the correlation event, as a function of the severity levels of all the root-cause node instances that are included in the range. From the list, select either **Average** or **Max.**

For example, if you select **Average**, the severity level of the correlation event is the average severity of all the root-cause instances.

12 Repeat steps 5 to 11 to define additional correlation events related to the selected node.

13 Click **OK** to save the settings you have defined. The new correlation rule is added to the Details dialog box.

The following information is displayed:

➤ **Description** – The description of the new correlation rule that is created.

➤ **Send Event on Category** – The category for which the event is sent.

14 Once you complete the affected node definitions, click **Finish**.

Editing a Correlation Rule

You can edit an existing correlation rule.

To edit a correlation rule:

1 In the Details dialog box, select the required correlation rule and click the **Edit** button.

2 Make the required changes (for details, see "Setting Multiple Conditions for Defining Affected Nodes" on page 287).

3 Click **Finish** to save the changes you have made.

Removing a Correlation Rule

You can remove an existing correlation rule.

To remove a correlation rule:

1 In the Details dialog box, select the correlation rule you want to delete and click **Delete**.

2 Click **Finish**.

Saving the Correlation Rule

After you have defined the correlation rule, the last step is to save it to the database. If you exit the Correlation Manager without saving, a confirmation window is displayed, requesting that you save your rule.

To save the correlation rule:

1 Display the correlation rule to be saved in the Topology map.

2 Click the **Save** button on the toolbar or select a different rule or tab in the View Explorer. A message is displayed, asking if you want to save the current correlation rule. The correlation rule you have created is saved to the database.

Note: If an active correlation rule is later deactivated (by clearing **Active** in the Properties dialog box), all events related to the rule are deleted from the system.

Importing a Correlation Rule

You can import XML files that contain saved correlation rules to your Correlation Manager. Use this option if you want to relocate correlation rules from one workstation to another.

Note: Before you import a correlation rule, you must verify that its attached TQL query appears in the **Attached TQL** list in the New dialog box. If the attached query does not exist in your query list, importing fails.

To import a correlation rule:

1 From the Correlation Manager, open the **Map** menu and click **Import** to open the Import dialog box. For details, see "Importing a TQL Query" on page 91.

2 Locate the correlation rule you want to import, select it and click **Import**. The imported correlation rule is added to the correlation rule list in the View Explorer.

Note: By default, an imported correlation rule is not active. To activate it, select it on the View Explorer and right-click. From the shortcut menu, select **Properties**. In the displayed Properties dialog box, select **Active**, and click **OK**.

Part VII

Enrichment Nodes and Rules

16

Introduction to Enrichment Nodes and Relationships

This chapter introduces the Enrichment Manager, which enables you to create enrichment nodes and rules and add them to a TQL query.

This chapter describes:	On page:
About Enrichment Nodes, Relationships, and Rules	295
Quick Tour of the Enrichment Manager	297
Understanding Enrichment Manager Modes	299
Working With Enrichment Manager Shortcut Menus	299
Enrichment Manager Tooltip	304
Working with the Enrichment Manager	304

About Enrichment Nodes, Relationships, and Rules

The Enrichment Manager provides you with greater flexibility in managing queries and designing views, and enables you to integrate prior knowledge of the infrastructure with data that is found and gathered through the Mercury Application Mapping discovery process.

Enrichment nodes and relationships differ from other Mercury Application Mapping CIs and relationships in that they are not actually discovered by discovery patterns. They are conceptual deductions that represent real relationships and CIs that cannot be discovered automatically by the discovery process.

Enrichment nodes and relationships are created as part of a TQL, whose other TQL nodes are regular ones, meaning TQL nodes that already exist in the database. The following example illustrates that a regular Network node is linked to a TOP2 enrichment node by an enrichment Member relationship.
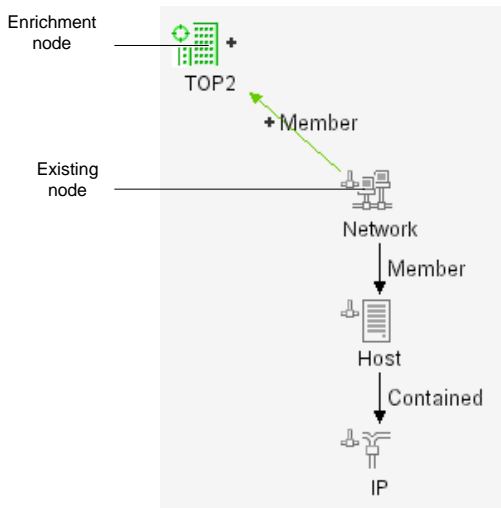
By placing an enrichment node within a TQL context, it receives data from the attributes of the other enrichment nodes in the TQL, and uses it to insert new information into the database that is inferred from an existing one, or represents existing information in a new way.

Since these enrichment nodes and relationships are usually characterized differently, their creation often entails a creation of new CITs (done through the CI Type Manager; for details, see "Introduction to the CI Type Manager" on page 377).

Enrichment CIs and Relationship Objectives

There are two main reasons for creating enrichment CIs and relationships:

➤ Enlarging the CMDB by adding enrichment CIs and relationships that are currently not included in it, but whose data is known or can be logically deduced from discovered CIs.

For example, suppose it is known that a certain relationship, such as **depend**, exists between A and B, and between B and C, as shown in the following figure:
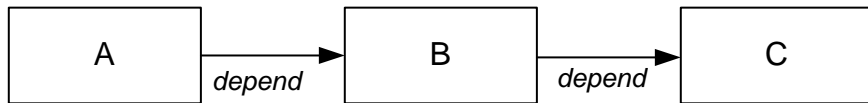
Based on this information, it is logical to deduce that between A and C a **depend** relationship also exists, as shown in the following figure:
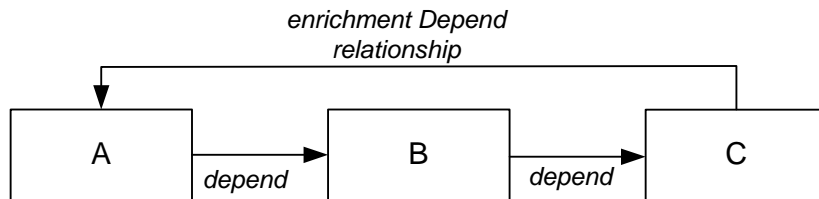
By using an enrichment rule, you can add the relationship between A and C to the CMDB.

➤ Creating informative representations of your infrastructure, by simplifying data representations that are based on query results of complex structures.

Quick Tour of the Enrichment Manager

The Enrichment Manager, which is displayed by selecting the **Enrichment Manager** tab in the View Explorer, enables you to create and define enrichment rules. These rules, which are based on a specific enrichment TQL query, can be used for several purposes:

➤ Creating new CIs and relationships for adding data to the CMDB

➤ Enabling additional representation options of existing data

➤ Updating the value of CI attributes in the database

➤ Deleting CIs from the CMDB by using TQL queries
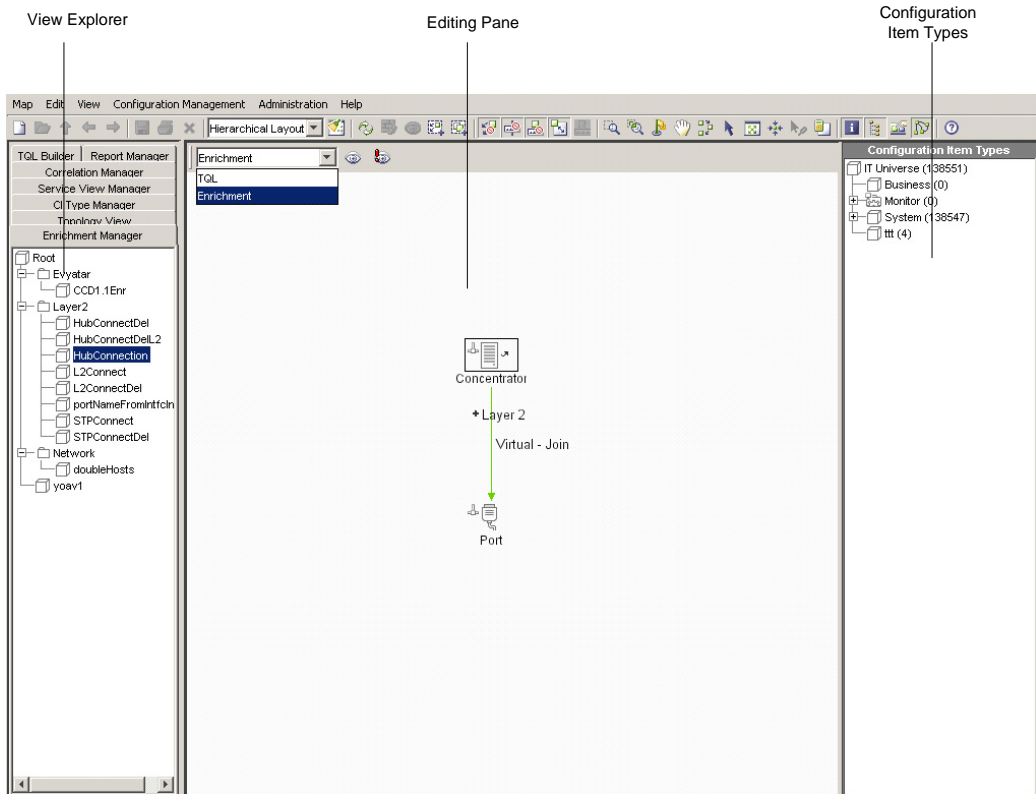
When the Enrichment Manager tab is selected, Mercury Application Mapping main window is divided as follows:

➤ **View Explorer –** Displays a hierarchical tree structure of the enrichment rules you have defined. If required, you can export enrichment rules and save them as XML scripts, which can be used for backup purposes.

Note: Each enrichment rule type is represented by a unique icon.

➤ **Editing Pane –** Displays the currently selected rule, which consists of TQL nodes that are defined in the TQL query and the relationships between them, and the enrichment nodes and relationships that are created and added to the rule.

➤ **Configuration Item Type Model –** Represents the CI Type Model and contains icons for each CIT as defined by the administrator. By clicking and dragging CITs to the editing pane, you can create new enrichment nodes. By defining the relationship between these enrichment nodes and existing nodes, you can create new enrichment relationships.

Understanding Enrichment Manager Modes

The Enrichment Manager has two modes in which you can work:
Enrichment mode and **TQL** mode.

Working in:

➤ TQL mode enables you to define an enrichment TQL query. For details, see "Using the TQL Node Wizard" on page 135.

➤ Enrichment mode enables to define enrichment relationships and nodes and add them to the TQL query you defined. For details, see "Creating Enrichment Nodes and Relationships" on page 310.

Working With Enrichment Manager Shortcut Menus

The shortcut menus that appear depend on which mode you are using.

➤ Working in Enrichment mode enables you to:

➤ Add enrichment nodes and relationships to the TQL query

➤ Update attribute definitions for enrichment nodes and relationships

➤ Delete enrichment nodes and relationships

Note: For details, see "Node/Relationship Shortcut Menu When Working in Enrichment Mode" on page 302.

➤ Working in TQL mode enables you to:

　➤ Create a TQL query using the TQL wizard

　➤ Add nodes and relationships to the query

　➤ Define node and relationship attribute conditions

　➤ Create a dependency graph

　➤ Display all of the instances found for each TQL node in a table

---

Note: For details, see "Node/Relationship Shortcut Menu When Working in TQL Mode" on page 303.

This section includes:

➤ "Enrichment Manager Enrichment Rule Shortcut Menu" on page 301

➤ "Node/Relationship Shortcut Menu When Working in Enrichment Mode" on page 302

➤ "Node/Relationship Shortcut Menu When Working in TQL Mode" on page 303

Enrichment Manager Enrichment Rule Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a selected enrichment rule in the View Explorer:

Option	Description
New	Enables you to define a new enrichment rule. For details, see "Creating an Enrichment Rule" on page 307.
New Folder	Enables you to create a new enrichment folder.
Save	Enables you to save the enrichment rule and node to the CMDB. (Enabled only when a new rule is created or when changes are made to an existing one.) For details, see "Saving the Enrichment Rule" on page 322.
Save As	Displays the Save As dialog box, enabling you to define a new name and description for the rule (used for creating a new rule based on an existing one). For details, see "Creating an Enrichment Rule Based on an Existing One" on page 321.
Delete	Enables you to delete the selected rule from the CMDB. For details, see "Deleting CIs From the CMDB Using Enrichment Rules" on page 324.
Export	Displays a standard Save As dialog box, enabling you to save the enrichment rule in XML format. This option can be used to move selected rules from one workstation to another, provided the related TQL query is also relocated. For details, see "Exporting Enrichment Rules" on page 321.
Properties	Enables you to change the properties of the rule (aside from the rule name and the Enrichment TQL query to which it is attached)."Updating CI Attributes Using Enrichment Rules" on page 313

Node/Relationship Shortcut Menu When Working in Enrichment Mode

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking an enrichment node in the editing pane when you are working in Enrichment mode. Enrichment relationships appear in green and enrichment nodes and relationships are displayed by an added indicator.

✦

Option	Description
Add Relationship	Enables you to add an enrichment relationship to the rule. Applicable to regular and enrichment nodes. For details, see "Adding TQL Nodes and Relationships to a Query" on page 61.
Delete	Enables you to delete the selected node/relationship. Applicable only to enrichment nodes and relationships. For details, see "Adding TQL Nodes and Relationships to a Query" on page 61.
Straighten	Enables you to straighten the relationship between two nodes. This option is only available for relationships that have angles. For details, see "Straightening a Relationship With an Angle" on page 86.
Update Relationship/Node	Enables you to update the attribute values of the selected node/relationships. This is applicable for both regular and enrichment nodes. For details, see"Updating CI Attributes Using Enrichment Rules" on page 313.
Delete Relationship/Node	Enables you to delete the selected node/relationship from the CMDB. Applicable to regular nodes and relationships only. For details, see"Deleting CIs From the CMDB Using Enrichment Rules" on page 324.
Clear	Deletes the attributes you defined for the selected node.
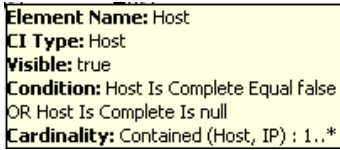
Node/Relationship Shortcut Menu When Working in TQL Mode

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking an enrichment node in the editing pane when you are working in TQL mode.

Option	Description
TQL Node Wizard	Builds a TQL query. For details, see "Using the TQL Node Wizard" on page 80.
Add Relationship	Displays the Add Relationship dialog box, enabling you to create a relationship for your TQL nodes by selecting it from a predefined list. For details, see "Adding Nodes and Relationships to a Query" on page 62.
TQL Node Definition	Displays the TQL Node Definition dialog box, enabling you to define the attribute conditions for the selected TQL node. For details, see "Setting TQL Node and Relationship Definitions" on page 68.
Delete	Enables you to delete the selected TQL node(s).
Add SubGraph	Creates a graph that represents additional TQL query data related to a specific node. For details, see "Creating a Dependency Graph" on page 86.
Show element's instances	Displays all of the instances found for each TQL node in a table. For details, see "Showing TQL Node Instances" on page 88.

Enrichment Manager Tooltip

Hold the pointer over a node or relationship to view its tooltip.

The tooltips contain the following information:

➤ **Element name** – the name of the node

➤ **CI Type** – The CI Type of the CI as defined in the Configuration Item Type Model.

➤ Definitions of the selected nodes and relationships. For example, the attribute conditions, as described in "Setting TQL Node and Relationship Definitions" on page 68.

Working with the Enrichment Manager

Use the Enrichment Manager to perform the following functions:

➤ Create enrichment TQL queries

➤ Create enrichment rules, as well as duplicate and delete existing rules

➤ Define enrichment nodes and relationships and add them to the CMDB

➤ Update the attribute values of existing CIs

➤ Delete CIs from the CMDB

➤ Export and import selected enrichment rules from XML scripts. This can be used to relocate enrichment rules from one workstation to another. (This can be performed providing the related TQL query is also relocated.)

➤ Zoom in and out of the editing pane to view the selected layer at different levels of magnification

➤ Select different layout views.

17

Defining Enrichment Nodes and Rules

This chapter explains how to define enrichment nodes and relationships and add them to an existing TQL query.

About Defining Enrichment Nodes and Rules

When creating an enrichment rule, you first name the rule and attach it to an existing enrichment TQL query or create a new one. You then proceed to create new enrichment nodes and relationships, by dragging and dropping the enrichment nodes from the CI Type Model to the scenario on the editing pane, and connecting them to existing TQL nodes by logical and virtual relationships. You then set the attribute values of the enrichment nodes and relationships for identifying the new CIs and relationships in the Mercury Application Mapping environment, providing them with meaningful information.

Enrichment Rule Workflow

Enrichment rules are created according to the following workflow:

➤ Create an enrichment rule. For details, see "Creating an Enrichment Rule" on page 307.

➤ Create enrichment nodes and relationships. For details, see "Creating Enrichment Nodes and Relationships" on page 310.

➤ Setting attribute definitions. For details, see "Updating CI Attributes Using Enrichment Rules" on page 313.

➤ Save the enrichment rule. For details, see "Saving the Enrichment Rule" on page 322.

Creating an Enrichment Rule

When creating a new enrichment rule, bear in mind that Mercury Application Mapping does not allow you make any changes in the TQL that are necessary for the rule. These changes include the following:

➤ Deleting, from the enrichment TQL, the nodes or relationships connected to the enrichment node, causing a node to become unattached to any relationship.

➤ You cannot change the cardinality definition to 0 (not required) of a node or relationship that is used in an enrichment rule. For example, when you define the relationship cardinality, and enter a value that defines the lower limits in an enrichment rule as 1 or more (required), you cannot change the cardinality definition to 0 (not required). For details on relationship cardinality, see "Setting TQL Node and Relationship Definitions" on page 124.

Note: The Enrichment TQL that serves as the basis for the enrichment rule, should comply with the validation restrictions. For details, see Chapter 5, "Validation Restrictions." If the TQL is not valid, it cannot be used for the creation of an enrichment rule.

To create an enrichment rule:

1 Click **New** on the toolbar to open the New dialog box.
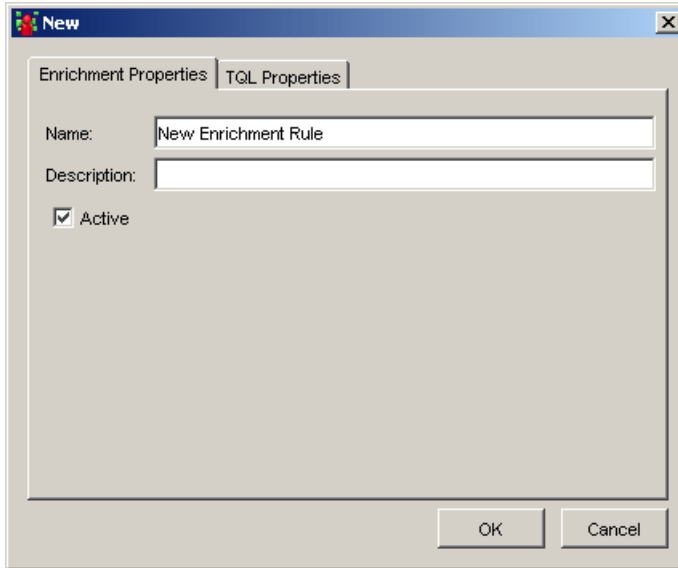
2 Click the **Enrichment Properties** tab.

3 In the **Name** box, enter a unique name for the enrichment rule.

4 (Optional) In the **Description** box, enter a description of the rule.

5 (Optional) Select the **Active** check box to activate the rule in the system as soon as it is saved.

Note: If you do not activate the rule now, you can do it later. Select the rule in the View Explorer and right-click the rule. From the shortcut menu, select **Properties**. In the Properties dialog box, select **Active**.

---

6 Click the **TQL Properties** tab to either create a new TQL or attach an existing one.

7 To create a new TQL, do the following:

> ➤ Select **Create New TQL**.

> ➤ In the **TQL Details** section, type a unique name for the TQL in the **Name** box.

> ➤ From the **Priority** list, select a priority level for the new TQL query (Low, Medium, High, Express). This setting determines how often the query is rerun automatically by the system to include updated information from the CMDB.

> ➤ (Optional) In the **Description** box, enter a description of the TQL query.

> ➤ If required, select **Is Persistent** to define whether you want the TQL to always remain in the system's memory. Use this option for TQLs that are used frequently rather than occasionally.

8 To attach a new TQL, do the following:

> ➤ Select **Attach to existing TQL.**

> ➤ From the **Name** list, choose the enrichment TQL you want to attach to the enrichment rule .

> ➤ From the **Priority** list, select a priority level for the new TQL query (Low, Medium, High, Express). This setting determines how often the query is rerun automatically by the system to include updated information from the CMDB.

> ➤ (Optional) In the **Description** box, enter a description of the TQL query.

> ➤ If required, select **Is Persistent** to define whether you want the TQL to always remain in the system's memory. Use this option for TQLs that are used frequently rather than occasionally.

9 Click **OK**. The enrichment rule is displayed in the View Explorer.

Creating Enrichment Nodes and Relationships

This section describes how to create enrichment nodes and to define the relationships between them and existing TQL nodes.

For a list of the relationships you can use to link two nodes and their definitions, see "Relationship Definitions" on page 543.

This section contains the following topics:

> ➤ "Adding Nodes and Relationships to Define a TQL Query" on page 311

> ➤ "Deleting a Node or Relationship" on page 312

Adding Nodes and Relationships to Define a TQL Query

This section describes how to add enrichment nodes and relationships to the enrichment TQL query.

To create enrichment nodes and relationships:

1 From the tree in the View Explorer, select the enrichment rule to which you want to add enrichment nodes and relationships.

2 At the top of the Enrichment Manager window, select **Enrichment** mode.

3 From the tree displayed in the Configuration Item Types pane, click and drag the Configuration Item Type you want to function as an enrichment node to the editing pane. Added enrichment nodes and relationships are displayed by an added indicator.

✛

Note: You can add more than one enrichment node to a rule.

---

4 The next step is to link the enrichment node to an existing TQL node or nodes, to provide the enrichment node with the needed context for its operation.

Note the following enrichment rule validations:

➤ You must link the new enrichment node to at least one of the existing TQL nodes in the rule.

➤ You cannot link enrichment nodes to one another.

➤ If the new enrichment node must be contained (according to its CIT definition) in another node, you must use the **Container link** relationship to connect this enrichment node to an existing TQL node.

➤ You cannot link an enrichment node to a TQL node that is not visible.

5 Select the enrichment node and the existing TQL node between which you want to create a relationship, by holding down CTRL and clicking.

6 When both items are selected, right-click to display the Add Relationship dialog box and select the relationship type that defines the connection between the two items.

Note: The relationships that appear in the Add Relationship dialog box are the ones that are defined in the CI Type Manager as the relationship between the two CITs. If you create a new CIT for the enrichment node, and the relationship you need does not appear in the Add Relationship dialog box, return to the CI Type Manager and add a relationship between the two CITs, as described in "Adding Relationships Between CITs" on page 400.

✚   Added enrichment relationships are displayed by an **Added** indicator.

The Add Relationship dialog box contains the following fields:

Field	Description
Node Label	The label of the selected nodes.
Relationship Direction	The direction of the relationship that indicates which node is dependent on the other.
Relationship Type	A valid relationship that defines the connection between the selected nodes.

7 Select the required direction of the relationship.

8 Click **OK**. The selected nodes are linked by the relationship you have selected.

Deleting a Node or Relationship

This section describes how to delete a node or relationship.

To delete a node or relationship:

Right-click the TQL node or relationship you want to delete and select **Delete**.

Using the TQL Node Wizard

You can use the TQL Node Wizard to build a TQL query. For details, see "Using the TQL Node Wizard" on page 80.

Showing TQL Node Instances

You can display all of the instances found for each TQL node in a table. For details, see "Showing TQL Node Instances" on page 88.

Updating CI Attributes Using Enrichment Rules

Use an enrichment rule to update the value of CI attributes in the CMDB, or to add data to attributes that currently do not have values (for details, see "Creating an Enrichment Rule" on page 307). You can use this option, for example, for adding a note simultaneously to all instances of a CI.

To update CI attributes, you open an enrichment TQL query in the Enrichment Manager, and update the attributes of one or more of its nodes.

This section contains the following topics:

➤ "Updating CI Attributes" on page 313
➤ "Enrichment Rule Validation" on page 314
➤ "Defining Node and Relationship Attribute Definitions" on page 314
➤ "Setting Attribute Definitions Using Simple and Complex Expressions" on page 316

Updating CI Attributes

This section describes how to update CI attributes.

To update CI attributes:

 1 At the top of the Enrichment Manager window, select **Enrichment** mode.

2 On the editing pane, right-click the node whose attribute(s) you want to update and select **Update Relationship/Configuration Item** to open the Node Definition dialog box.

3 Select the attribute you want to update and click **Edit**. Depending on the attribute type you selected, one of two Set Attributes dialog boxes is displayed.

4 Enter the updated value as described in "Defining Node and Relationship Attribute Definitions" on page 314 and "Setting Attribute Definitions Using Simple and Complex Expressions" on page 316. Updated CIs are displayed by an arrow indicator.

5 Click **OK** to save your changes.

Note: If you want to cancel the update and you have not yet saved the rule, right-click the updated CI and select **Clear**. If you already saved the rule, you cannot reverse the update.

---

6 To save the rule and update the CIs, click **Save** on the toolbar.

Enrichment Rule Validation

You must fill in the value of the key attribute(s) of the enrichment node. The method you use to define these values determines the number of instances created:

➤ If you enter a dynamic value, you can create numerous instances. For example, entering the dynamic value of a host_key for for the CIT **Host** or an IP address attribute for the CIT **IP**.

Defining Node and Relationship Attribute Definitions

This section describes how to define node and relationship attributes for all attributes other than the type **String**.

To define node and relationship definitions for all attributes other than the type String:

1 At the top of the Enrichment Manager window, select **Enrichment** mode.

2 On the editing pane, right-click the node whose attribute(s) you want to update and select **Update Relationship/Configuration Item** to open the Node Definition dialog box.

3 Select the attribute you want to update and click **Edit**. The following Set Attributes dialog box is displayed.
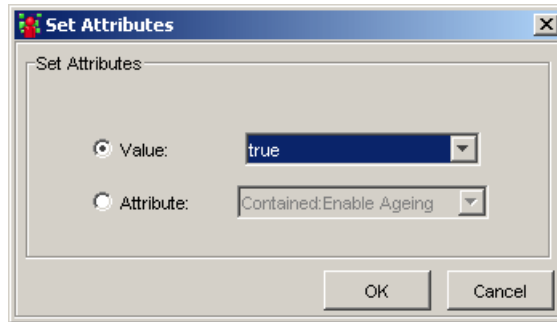
4 To define a constant value, do the following:

➤ Select **Value**.

➤ Enter the required value in the **Value** box or select a value from the **Value** list.

5 To define a variable, do the following:

➤ Select **Attribute**.

➤ Select a value from the **Attribute** list.

➤ Click **OK** to save the changes.

The attribute definition appears in the **Value** column in the Node Definition dialog box.

Setting Attribute Definitions Using Simple and Complex Expressions

This section describes how to define node and relationship attributes of the type **String** using simple or complex expressions. You can also create complex expressions using regular expression format.

To define node and relationship definitions:

1 At the top of the Enrichment Manager window, select **Enrichment** mode.

2 On the editing pane, right-click the node whose attribute(s) you want to update and select **Update Relationship/Configuration Item** to open the Node Definition dialog box.

3 Select the attribute you want to update and click **Edit**. The following Set Attributes dialog box is displayed.
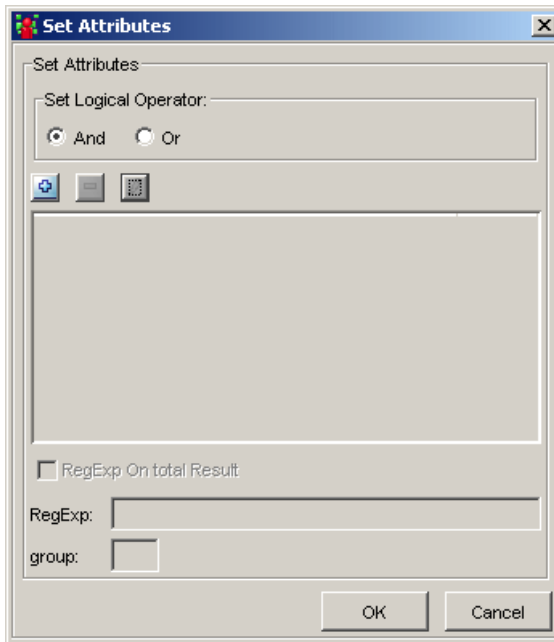
Note: This dialog box only appears when you select an attribute of the type **String**.

4 Click the **Add Object** button to add an attribute value in the Type dialog box.

5 To define a constant value using a simple expression, select **Simple**.

6 Enter the required expression in the **Simple** box.

7 To define an attribute using attributes from other nodes in the TQL, select **Complex**.

8 From the **Complex** list, select the required attribute.

9 If required, you can also select **RegExp on Total Result** to add a regular expression using regular expression syntax.

10 In the **RegExp** box, enter the regular expression pattern. This is the structure of the selected attribute. For examples of how to use regular expression syntax, see "Using Regular Expressions" on page 398.

11 In the **Group** field, enter the group number. This is the part of the regular expression pattern to focus on when creating the attribute. A set of parentheses () constitutes a group.

12 Click **OK** to save your attribute definitions. The attribute definition appears in the Set Attributes dialog box.

13 To define additional attribute values using simple or complex expressions, repeat steps 4 to 12.

14 In the **Select Logical Operator** area, select either **And** or **Or** to link multiple attribute definitions.

➤ Select **And** to use all the attribute values listed in the Set dialog box.

➤ Select **Or** to use the first attribute value listed in the Set dialog box that is not empty.

Following is an example of an IP address created by using simple and regular expressions.
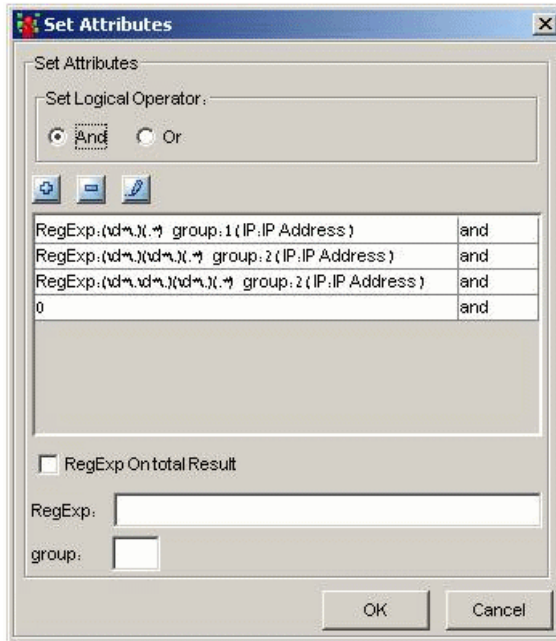
15 To remove an attribute value, select the attribute value you want to delete and click the **Remove Object** button.
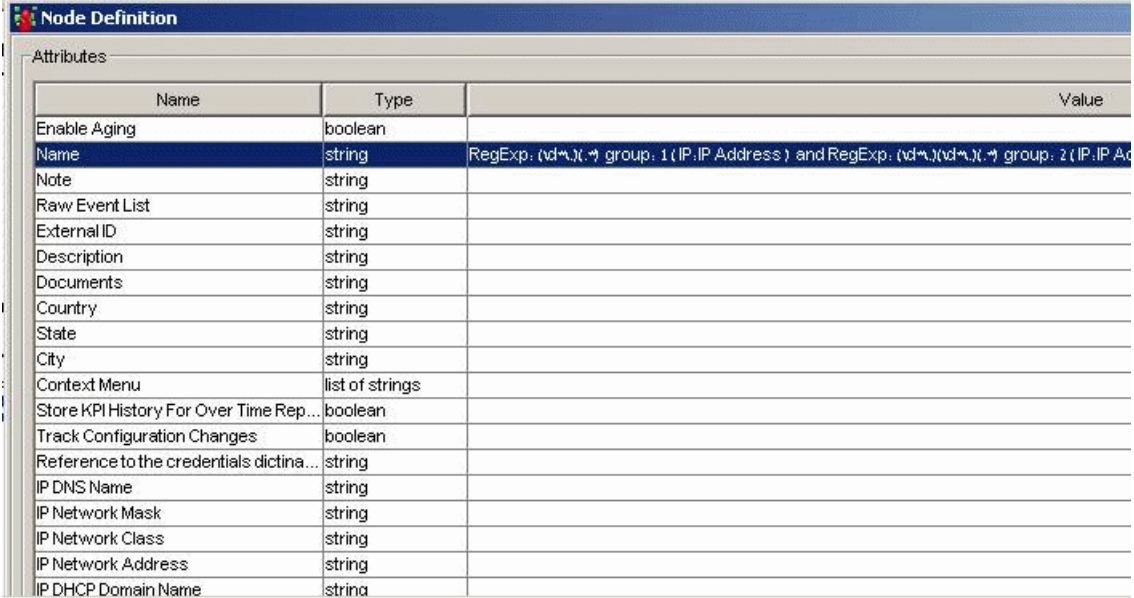
16 To edit an attribute value, select the attribute value you want to edit and do the following:

➤ Click the **Edit Object** button to open the Set Attributes dialog box.

➤ Edit the attribute value as described in steps 5 through 11.

17 Click **OK**. The attribute definition appears in the **Value** column in the Node Definition dialog box.

Showing Enrichment Results

Mercury Application Mapping describes how to calculate the number of instances that were created by the enrichment rule and how to delete them from the CMDB.

This section has the following topics:

➤ "Calculating the Number of Instances That Were Created from an Enrichment Rule" on page 320

➤ "Removing the Instances an Enrichment Rule Created From the CMDB" on page 320

Calculating the Number of Instances That Were Created from an Enrichment Rule

This section describes how to calculate the number of Instances that were created as a result of an Enrichment Rule.

To calculate the number of instances that were created from the enrichment rule you defined:

1 In the View Explorer, select the required enrichment rule.

2 Click the **Created Objects result count** icon. The number of TQL node instances and relationships that were created appear next to the enrichment nodes/relationships.
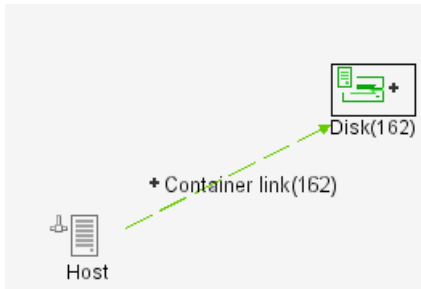
Removing the Instances an Enrichment Rule Created From the CMDB

This section describes how to remove the instances created from an enrichment rule from the CMDB.

To remove the instances that an enrichment rule created:

1 In the View Explorer, select the required enrichment rule.

2 Click the **Remove enrichment results** icon. The instances are removed from the CMDB.

Creating an Enrichment Rule Based on an Existing One

You can create an enrichment rule based on an existing one.

To create an enrichment rule based on an existing one:

1 In the View Explorer, right-click the rule on which you want to base a new one and select **Save As**. The Save As dialog box opens.

2 Click the **Enrichment Properties** tab.

3 In the **Name** box, enter a unique name for the enrichment rule.

4 (Optional) In the **Description** box, enter a description of the rule.

5 (Optional) To activate the rule in the system as soon as it is saved, select **Active**.

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

Exporting Enrichment Rules

You can save the enrichment rule in XML format. This option can be used to move selected rules from one client to another, provided the related TQL query is also relocated.

To export an enrichment rule:

1 Right-click the enrichment rule whose XML file you want to export and click **Export**. The Save as dialog box opens.

2 Browse to the location where you want to save the enrichment rule's XML file and click **Save**.

Saving the Enrichment Rule

After you have defined the enrichment rule, the last step is to save it to the database. If you exit the Enrichment Builder without saving, a confirmation message is displayed requesting that you save your rule.

To save the enrichment rule:

Select the rule you want to save and click **Save** on the toolbar or select a different rule or tab in the View Explorer. A message is displayed, asking if you want to save the current enrichment rule. Click **OK** to confirm the saving. The enrichment rule you have created is saved to the database.

Note: If an error message that prevents you from saving the rule is displayed, check whether you followed the rule validations. For details, see "Enrichment Rule Validation" on page 314.

Importing an Enrichment Rule

You can import XML files that contain saved enrichment rules to the Enrichment Manager. Use this option if you want to relocate enrichment rules from one workstation to another.

Note: Before you import an enrichment rule, you must verify that its attached TQL query appears in the **Attached TQL** list on the New dialog box. For details, see "Creating an Enrichment Rule" on page 307. If the attached query does not exist in your query list, importing fails.

To import an enrichment rule:

1 From the Enrichment Manager, open the **Map** menu and click **Import** to open the Import dialog box.

2 Locate the enrichment rule you want to import, select it, and click **Import**. The imported enrichment rule is added to the enrichment rule list in the View Explorer.

Activating an Imported Enrichment Rule

By default, an imported enrichment rule is not active. This section describes how to activate an enrichment rule that has been imported.

To activate an imported enrichment rule:

1 Right-click the rule you want to activate in the View Explorer and select **Properties** to open the Properties dialog box.

2 In the **Enrichment Properties** tab, select **Active**.

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

Deleting CIs From the CMDB Using Enrichment Rules

The Enrichment Manager enables you to delete CIs from the CMDB by using TQL queries. You create a specific enrichment rule that is designed to delete a CI. You can use this option, for example, for removing unnecessary data from the CMDB, or for automatically deleting copies of, or partial, CIs.

To delete a CI from the CMDB:

1 On the editing pane, right-click the node whose instances you want to delete from the CMDB and select **Delete Node/Relationship**. Deleted CIs are displayed by a deleted indicator.

Note: To cancel the deletion if you have not yet saved the rule, right-click the deleted CI and select **Clear**. If you have saved the rule, you cannot reverse the deletion.

---

2 To save the rule and update the CIs, click **Save** on the toolbar.

18

Creating a Sample Enrichment Rule

This chapter describes how to create a sample Enrichment rule and then display the results in the Topology View.

Creating an Enrichment Rule

This chapter leads you step-by-step through the process of creating a sample enrichment rule. The purpose of this exercise is to create an enrichment rule where, for every instance in the CMDB in which a **Port** CI and **Interface Index** CI are connected by a **Join** relationship, the name of the **Port** CI is created\updated from either the interface index name or the interface index description of the **Interface Index** CI.

To create a TQL query:

 1 In the Enrichment Manager, click the **New** button to open the New dialog box.

 2 Click the **Enrichment Properties** tab.
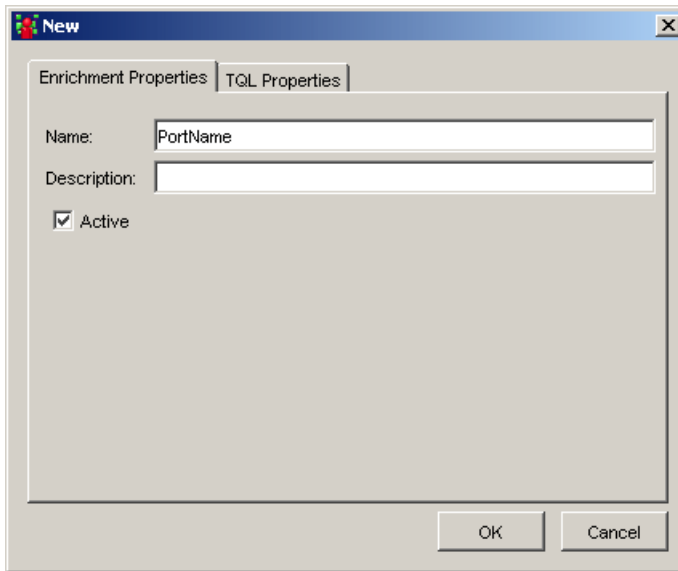
 3 In the **Name** box, enter **PortName**.

4 Click the **TQL Properties** tab.
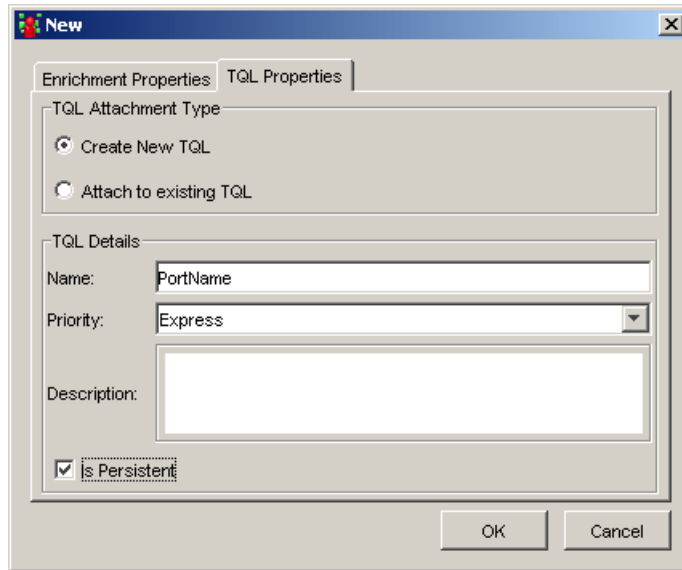
5 Select **Create New TQL**.

6 Select **Is Persistent** to ensure that the TQL always remains in the system's memory.

7 Click **OK**. The **PortName** enrichment rule is displayed in the View Explorer.
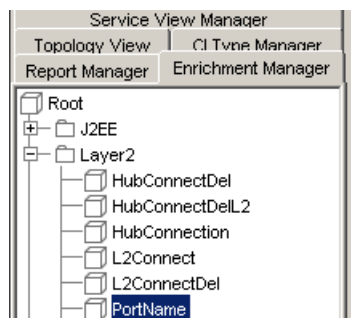
Adding Nodes to the Port Name Enrichment Rule

This section describes how to add nodes and relationships to the **PortName** enrichment rule.

To add nodes and relationships to the PortName enrichment rule:

 1 At the top of the Enrichment Manager window, select **TQL** mode.

 2 With the **PortName** enrichment rule selected, click and drag the following TQL nodes to the editing pane from the Configuration Item Type Model on the right:

   ➤ **Physical Port**

   ➤ **Interface Index**

 3 Select the **Physical Port** and **Interface Index** nodes and right-click.

 4 Choose **Add Relationship** to open the Add Relationships dialog box.

 5 Select the required direction of the relationship.

 6 Click **Advanced**.

 7 Open the **Relationship** list and select **Function Relationship**.

 8 Select **Virtual - Join link**.

 9 In the **Join Definition** area, click the **Add an attribute definition** button to open the Conditions dialog box.

 10 From the **Physical Port Attribute** list, choose **Port Interface Index**.

 11 From the **Operator** list, choose **Equal**.

 12 From the **Interface Index Attribute** list, choose **Interface Index**.

 13 Click **OK**.

 14 Click the **Add an attribute definition** button to open the Conditions dialog box.

 15 From the **Physical Port Attribute** list, choose **Container**.

 16 From the **Operator** list, choose **Equal**.

 17 From the **Interface Index Attribute** list, choose **Container**.

 18 Click **OK**.

19 Click **Finish**. The **Physical Port** and **Interface Index** nodes are now connected by a **Virtual - Join** relationship.

The TQL now looks like this:
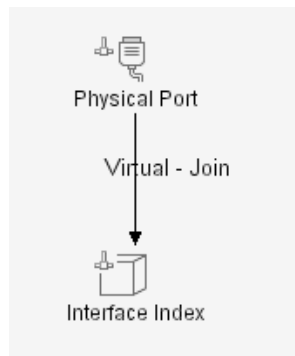
Note: For details on TQL validation restrictions, see Chapter 5, "Validation Restrictions."

20 To save the **PortName** TQL, click the **Save** button on the Enrichment Builder toolbar.

Defining the Enrichment Rule

After creating a TQL and creating a relationship between the **Physical Port** and **Interface Index** nodes, you update the **Physical Port** node using an enrichment definition.

To update the Physical Port node using an enrichment definition:

1 In the View Explorer, select the **Port Name** enrichment rule.

2 At the top of the Enrichment Manager window, select **Enrichment** mode.

3 Right-click the **Port** node and select **Update Relationship/Node Item** to open the Node Definition Dialog box.

4 Select the **Name** attribute and click **Edit** to open this Set Attributes dialog box..
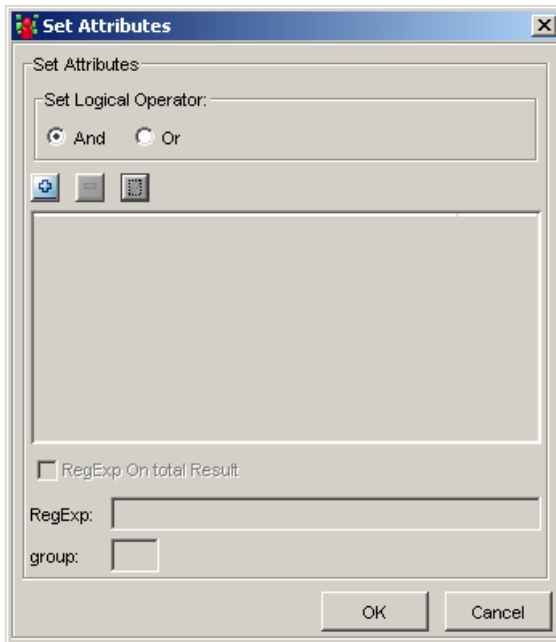
5 Click the **Add Object** button to open the Type dialog box.
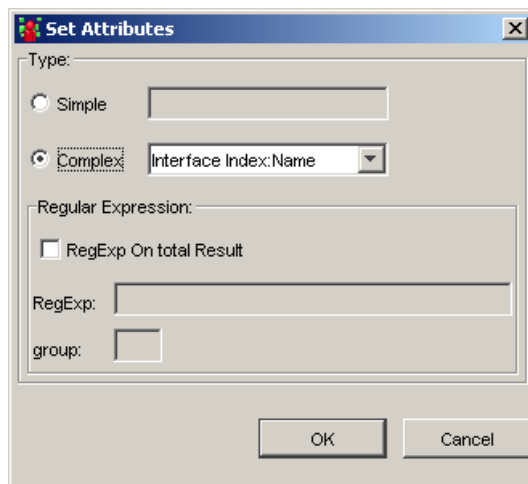
6 Select **Complex**.

7 From the **Complex** list, choose **Interface Index Name**.

8 Click **OK** to save your changes.

9 Click the **Add Object** button to open the Type dialog box.

10 Select **Complex**.

11 From the **Complex** list, choose **Interface Index:Interface Index Description**.

12 Click **OK** to save your changes.

13 In the **Set Logical Operator** area, select **Or**.

The Set Attributes dialog box now looks like this.
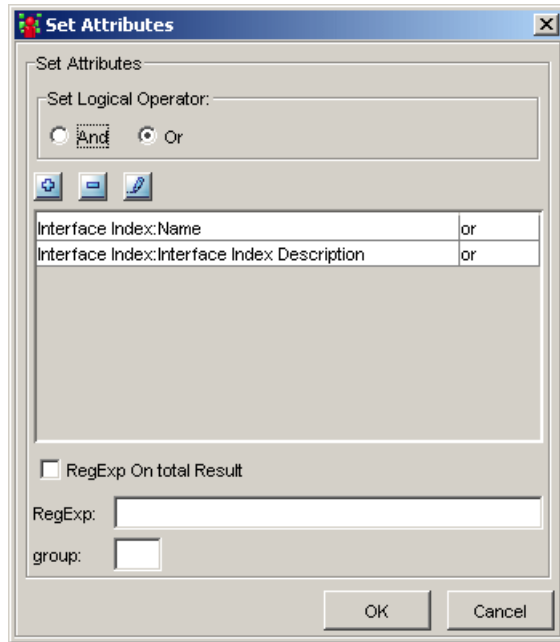
14 Click **OK**. The **Value** field for the **Name** attribute has been updated as follows:

Interface Index: Name or Interface Index:Interface Index Description
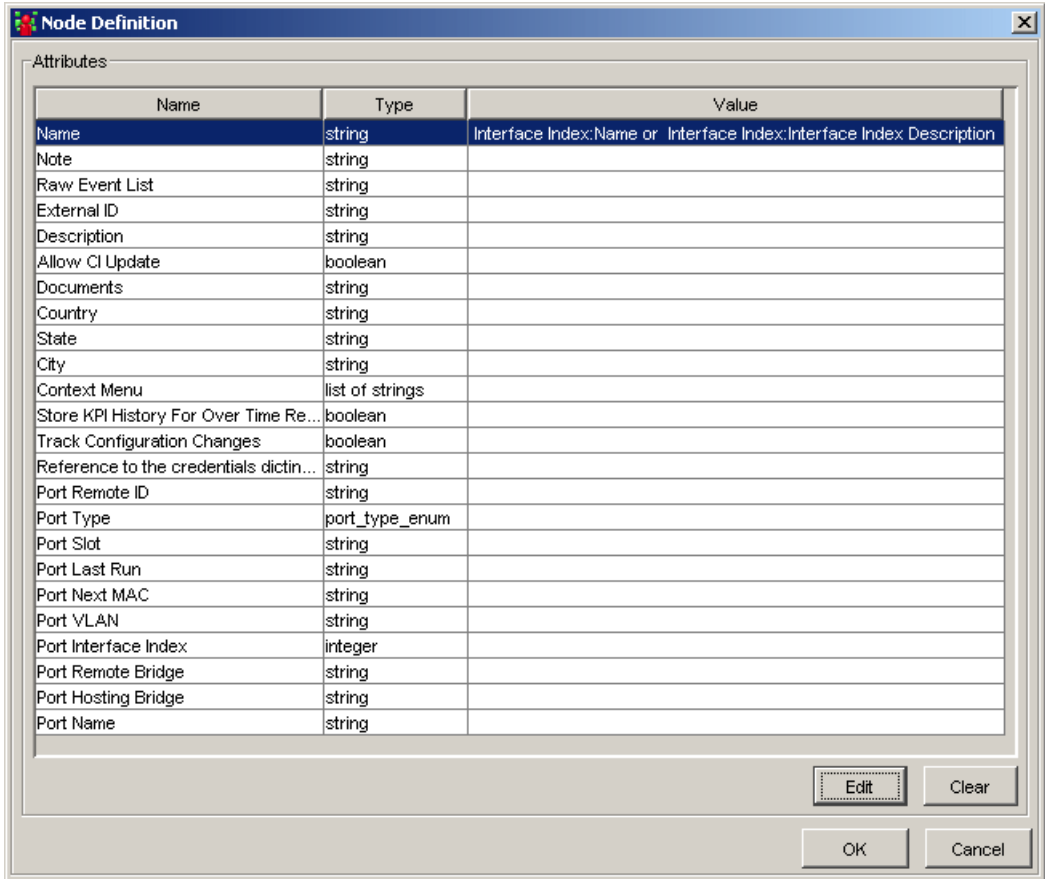
15 Click **OK** to save your changes.

The **Physical Port** node now has an arrow indicator to indicate that it has been updated with an enrichment definition.
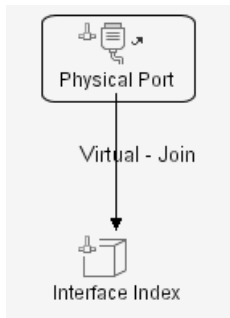
Creating a New View

This section describes how to display the enrichment result.

To create a new view:

1 In the Service View Manager, click the **New** button to open the New dialog box.

2 In the Create New View dialog box, enter the following information:

➤ In the **View Name** box, enter **Port_Name**.

➤ Select **Merge Identical Instances**.

3 Click the **TQL Properties** tab.

4 Select **Create New TQL**.

5 Click **OK** to save your definitions and close the Create New View dialog box. The **Port_Name** view appears in the View Explorer.

6 With the **Port_Name** view selected, click and drag the following TQL nodes on to the editing pane:

➤ **Physical Port**

➤ **Interface Index**

These are the TQL nodes that are included in the query.

333

7 Select the **Physical Port** and **Interface Index** nodes and right-click.

8 Choose **Add Relationship** to open the Add Relationships dialog box.

9 Select the required direction of the relationship.

10 Click **Advanced**.

11 Open the **Relationship** list and select **Function Relationship**.

12 Select **Virtual - Join link**.

13 In the **Join Definition** area, click the **Add an attribute definition** button to open the Conditions dialog box.

14 From the **Physical Port Attribute** list, choose **Port Interface Index**.

15 From the **Operator** list, choose **Equal**.

16 From the **Interface Index Attribute** list, choose **Interface Index**.

17 Click **OK**.

18 Click the **Add an attribute definition** button to open the Conditions dialog box.

19 From the **Physical Port Attribute** list, choose **Container**.

20 From the **Operator** list, choose **Equal**.

21 From the **Interface Index Attribute** list, choose **Container**.

22 Click **OK**.

23 Click **Finish**. The **Physical Port** and **Interface Index** nodes are now connected by a **Virtual - Join** relationship.

24 To save the **Port_Name** view, click the **Save** button on the toolbar.

25 Verify that a checkmark is displayed beside the **Port_Name** view in the View Explorer. This indicates that the view is to be displayed in the Topology View. If the checkmark is not displayed, right-click the view in the View Explorer and select **Add/Remove to Topology View**.

Displaying the Enrichment Results

This section describes how to view the results of the enrichment rule you created.

To view the results of the enrichment rule you created:

1 Click the **Topology View** tab.

2 In the View Explorer, select the **Port_Name** view.
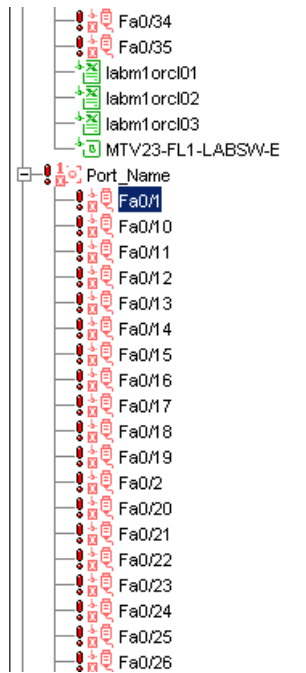
Your view should look similar to the following (depending on your system's structure):
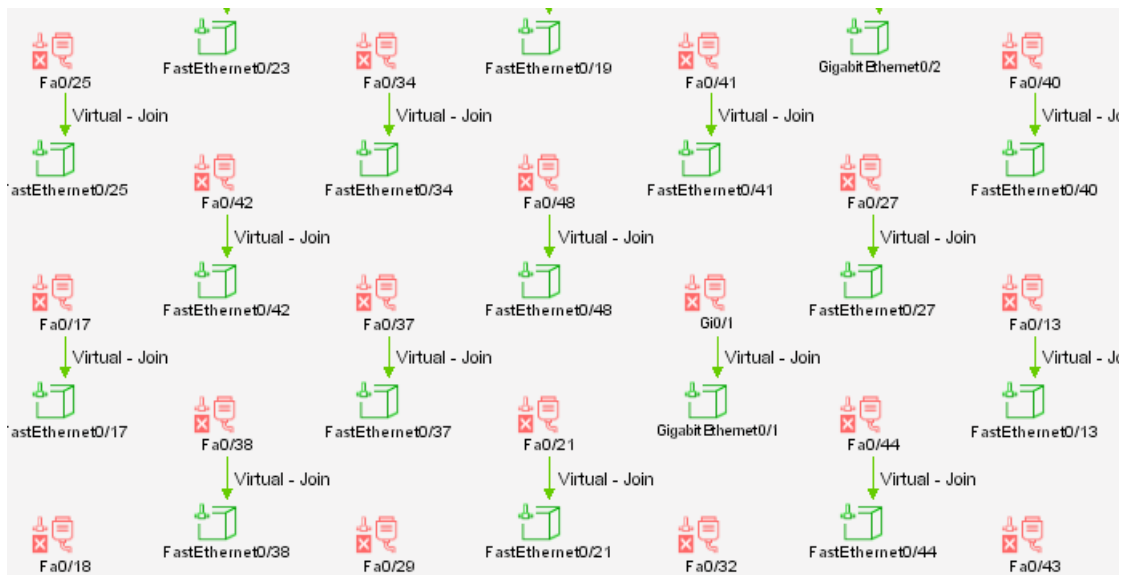
In this case, the name of the **Port** CI, **Fa0/41**, was taken from the **Name** attribute definition of the **Interface Index** CI.

To verify this, right-click the **Interface Index** CI and choose **Show CI Attributes**. The following illustration shows that the **Name** attribute definition is **Fa0/41**.
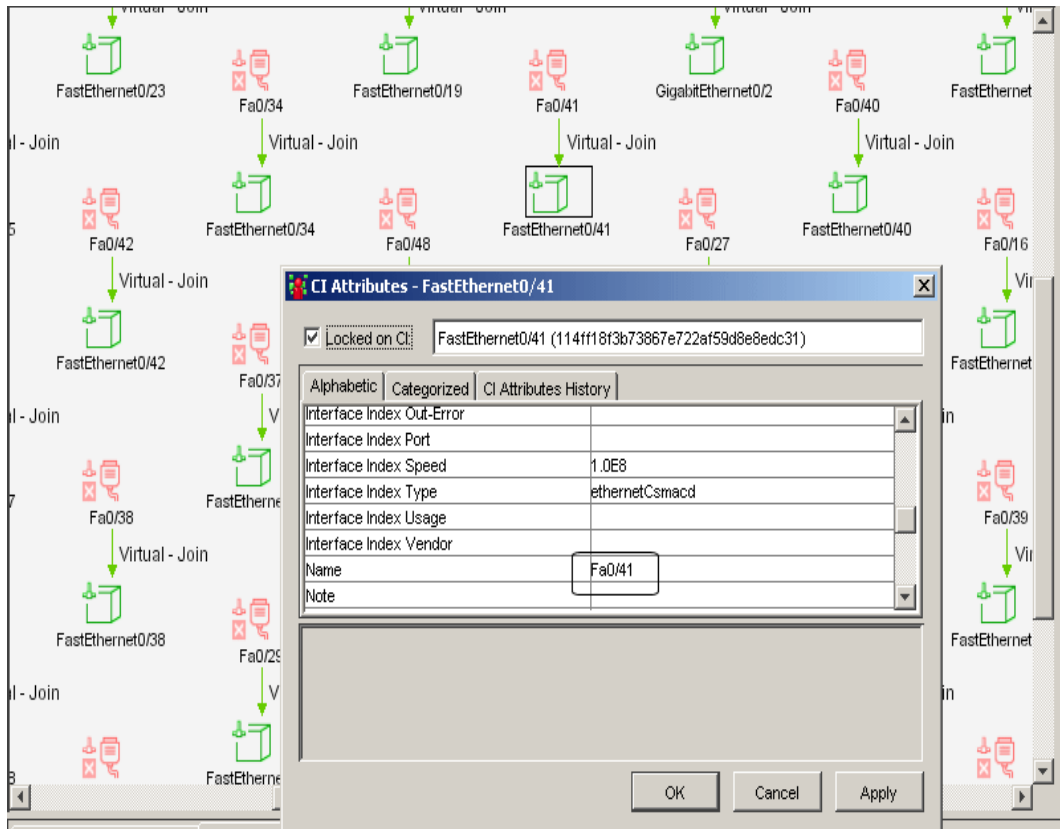
Part VIII

System Reports

19

Introduction to the Report Manager

This chapter introduces the Report Manager, which enables you to create reports.

This chapter describes:	On page:
About the Report Manager	342
Report Manager Shortcut Menus	343
Working with the Report Manager	345

About the Report Manager

The Report Manager is displayed by selecting the Report Manager tab in the View Explorer. It enables you to define system reports about selected Report nodes in the managed world. Each report is based on a specific Report TQL query, which is created especially for it. (If required, you can define multiple reports for each query.) The Report Manager enables you to define exactly what information appears in each report, including the order in which Report nodes are listed and their placement relative to other Report nodes, how the information within each Report node is sorted and the maximum number of rows of information to include for each Report node.

System reports can be displayed in three formats (HTML, Table, and Excel), and can either be viewed immediately or saved for later use.
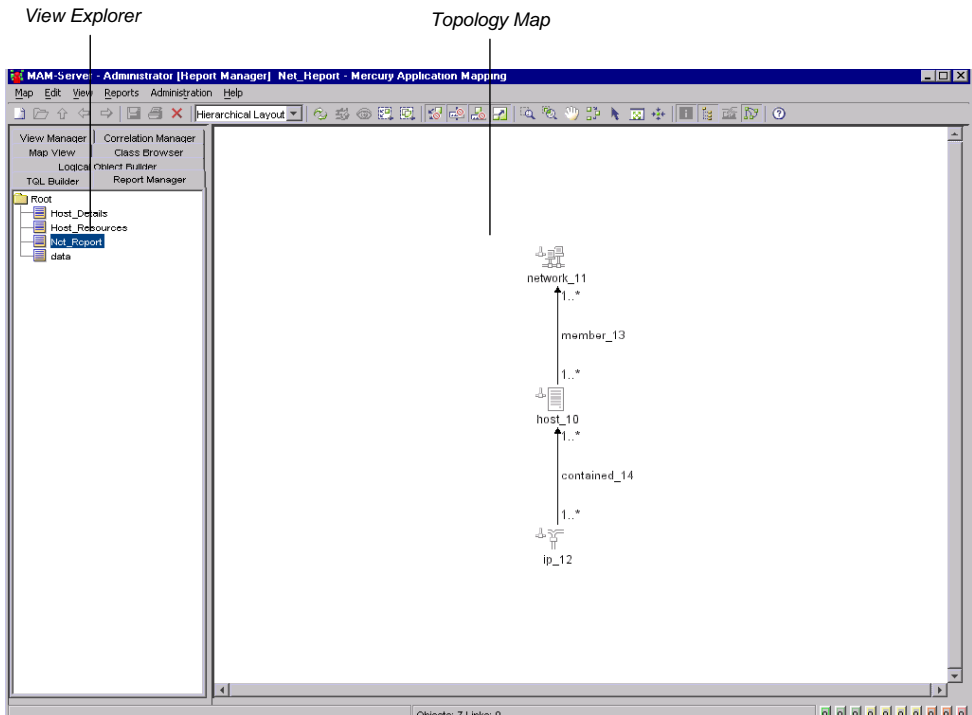
When the Report Manager tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **View Explorer –** Displays a hierarchical tree structure of the reports you have defined, each of them can be used only with the query to which it is attached. If required, you can export reports and save them as XML scripts, which can be used for backup purposes.

➤ **Topology Map –** Displays the currently selected report, which consists of Report nodes and the relationships between them. The displayed scenario is defined as a Report TQL query in the Enrichment Manager.

Report Manager Shortcut Menus

The Report Manager contains different shortcut menus, depending on your selection, as follows:

➤ "Report Shortcut Menu" on page 343

➤ "Report Node Shortcut Menu" on page 344

➤ "Relationship Shortcut Menu" on page 345

Report Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a selected report in the View Explorer:

Option	Description
New	Enables you to define a new System Report.
New Folder	Enables you to create a new report folder.
Save	Enables you to save the report to the database. (Enabled only when a report is created or when changes are made to an existing one.)
Save As	Displays the Save As dialog box, enabling you to define a new name and description for the report (used for creating a new report based on an existing one).

Option	Description
Delete	Enables you to delete the selected report from the database.
Generate	Enables you to generate the report, and display it or save it for future use.
Export	Displays a standard Save As dialog box, enabling you to save the report as an XML script. This option can be used to move selected reports from one workstation to another, provided the related TQL query is also relocated.
Properties	Enables you to change the properties of the report (apart from the report name and the Report TQL query to which it is attached).

Report Node Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a Report node in the topology map:

Option	Description
Report Node Definition	Displays the Report Node Definition dialog box, enabling you to determine the content of the report.
Clear Node Definition	Erases the report definition for the selected Report node.
Node Order	Displays the Set Node Order dialog box, enabling you to determine the order in which Report node information is displayed in the report, that is, the order of the report's columns.

Relationship Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking a relationship in the topology map:

Option	Description
Sibling	Defines one Report node as a sibling of the Report node to which it is linked. In the report, the two Report nodes appear side by side.
Child	Defines one Report node as a child of the Report node to which it is linked. In the report, the child Report node appears beneath (and indented from) the parent Report node.
Node Order	Displays the Set Node Order dialog box, enabling you to determine the order in which Report node information is displayed in the report.

Working with the Report Manager

Use the Report Manager to perform the following functions:

➤ Create new reports, as well as duplicate and delete existing reports

➤ Define the information to be included in the report, including which Report nodes should appear in the report, the parameter reported for each Report node and the function to be performed on the information

➤ Define the sort order of the information reported for each Report node

➤ Determine the order of Report nodes in the report

➤ Define which Report nodes appear as children beneath parent Report nodes

➤ Export and import selected views to and from XML scripts. This can be used to relocate reports from one workstation to another. (This can be performed providing the related Report TQL query is also relocated.)

➤ Zoom in and out of the topology map area to view the selected layer at different levels of magnification

➤ Select different layout views

20

Creating System Report Templates

This chapter explains how to create system report templates.

About Creating System Report Templates

When creating a System report template, the first step is to name the report and attach it to a specific report TQL query. You then proceed to define layout and information options for each of the nodes to be included in the report. This includes defining which parameter(s) are included for each node and the function (such as sum or max) that is performed on them. You must also define how the information should be sorted and the order and placement of each node in the report. When the report is complete, you select the format in which it is to be displayed and saved, and can either view it right away or save it for later viewing.

Report Manager Workflow

Reports are created according to the following workflow:

➤ Define a report (for details, see "Creating a Report" on page 349).

➤ Define the report nodes (for details, see "Defining the Report Nodes" on page 349).

➤ Set the node order (for details, see "Setting the Node Order" on page 354).

➤ Refine the report layout (for details, see "Refining the Report Layout" on page 355).

➤ Save the report template (for details, see "Saving the System Report Template" on page 356).

➤ View and save the generated report (for details, see "Viewing and Saving the Generated Report" on page 357).

Creating a Report

You name the report and attach it to a specific Report TQL. You also define presentation options, such as the report title and subtitle.

To define a report:

1 Select the Report Manager tab in the View Explorer pane.

2 Click the **New** button on the toolbar or open the **Map** menu and click **New** to open the Report Definition dialog box.

3 In the **Report Name** box, enter a unique name for the report. This name later appears in the system report list and helps you identify the report.

The name must be one word.

4 (Optional) In the **Report Description** box, enter a description of the report.

5 From the **Attached TQL** list, select the Report TQL on which you want to base the report.

6 In the **Presentation** area, define report presentation options, as follows:

➤ In the **Report Title** box, enter the title of the report. This title appears at the head of the report.

➤ In the **Report Sub-Title** box, enter the sub-title of the report. This title appears underneath the title at the head of the report.

7 Click **OK**. The new report is displayed in the View Explorer and in the Topology Map panes.

Defining the Report Nodes

The next step in the procedure is to define the nodes that are included in the report, that is, the report's columns. This includes determining:

➤ The information to be displayed for each node (attributes and functions)

➤ The maximum number of rows to allow for each node

➤ The node label (that is, the column title)

➤ The order of the information displayed for each node

After completing the report definition for one node, repeat the entire process for the other nodes that are to be included in the report.

Only defined nodes are included in the final report. However, other nodes are included in the report calculations, since they are part of the TQL that the report reflects, and they can influence the data that is displayed in it. For example, in a Report TQL, which consists of a linked host and IP, the IP has an attribute condition that states that only IPs that are in a critical state are included in the query results. Consequently, the report includes only hosts that are connected to IPs in a critical state, even if the IPs themselves are not included in the report.

To define the report nodes:

Note: The nodes to be included in the reports must be connected to each other through the TQL scenario. You cannot define nodes that are not linked to at least one other node that is included in the report.

---

1 Right-click a node in the scenario displayed in the Topology map, and select **Report Node Definition** to open the Report Node Definition dialog box.

2 In the **Title** box, select the appropriate title.

3 In the **Max Row Number** box, enter the maximum number of rows related to this node to include in the report. For example, if the selected node is an IP, you might set a maximum of 30 IPs to include in the report.

4 Click the **Add** button to display the Column Definition wizard.

You define the node's columns that consist of the node attributes and functions.

5 Select **Attribute values columns** and click **Next**.

6 To add a column to the report relating to the selected node:

➤ In the **Attribute List** area, select an attribute and click the single left-to-right arrow. The selected attribute moves to the **Column List** area.

➤ To remove an attribute from the **Column List** area, select it and click the single right-to-left arrow.

➤ To include all attributes of the node in the report, click the double left-to-right arrow.

➤ To remove all attributes from the **Column List** area, click the double right-to-left arrow.

➤ By default, column titles take the same name as the selected attribute. To change the title, select the attribute in the **Column List** area, change the title in the **Column Title** box, and click **Apply**.

➤ Add additional attributes, if required.

➤ Click **Finish** to return to the Report Node Definition dialog box.

7 Click the **Add** button to display the Column Definition Wizard.

➤ Select **Function**, and select the function to use to calculate the results on these columns.

Note: If you are using one of the following functions for a certain node—**Average**, **Count**, **Min**, or **Max**—you cannot display another node after it in the report.

---

8 (Optional) If the selected function requires additional definitions, click **Next** to display a wizard and define the necessary parameters.

9 Click **Finish** to close the Column Definition Wizard. The new column definitions are displayed in the Column Title tab of the Report Node Definition dialog box, as shown in the following figure:
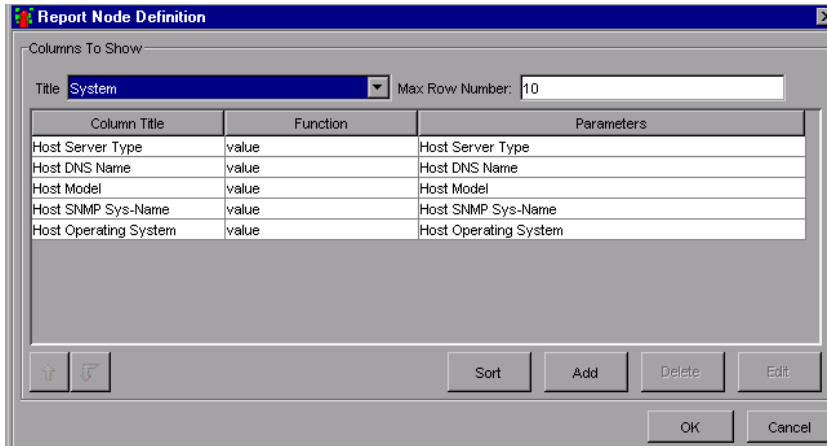
Each row in the table represents a column that appears in the report.

10 Repeat steps 4 to 6 to add additional column definitions.

Note:

➤ You can change the order of the columns (for details, see "Setting the Node Order" on page 354).

➤ You can sort the information in the columns (for details, see "Sorting Column Information" on page 353).

---

11 Click **OK** to close the Report Node Definition dialog box.

12 Repeat this procedure for the other nodes in the Report TQL scenario that are to be included in the report.

To edit an attribute in the report:

1 Right-click a node in the scenario displayed in the topology map, and select **Report Node Definition** to open the Report Node Definition dialog box.

2 To edit an attribute, select the row and click **Edit** (see above for more details).

3 Click **OK**.

To delete an attribute in the report:

1 Right-click a node in the scenario displayed in the topology map, and select **Report Node Definition** to open the Report Node Definition dialog box.

2 To delete an attribute, select the row and click **Delete**.

3 Click **OK**.

To clear a node's definition:

Right-click the required node and select **Clear Node Definition**.

Sorting Column Information

You can sort the column information in the system report.

To sort a column information:

1 Right-click a node in the scenario displayed in the Topology map, and select **Report Node Definition** to open the Report Node Definition dialog box.

2 Select the appropriate column and click **Sort** to open the Sort dialog box.

3 Sort the column information, as follows:

➤ From the **Sort Items by** list, select the primary sort column, and choose whether the information is to be sorted in ascending or descending order.

➤ (Optional) Define up to three secondary sort columns from the **Then by** lists. For each secondary sort, select **Ascending** or **Descending** to sort the information in ascending or descending order.

➤ Click **OK** to close the Sort dialog box and to re-display the Report Node Definition dialog box.

Setting the Node Order

After defining each of the nodes to include in the report, you can determine the order in which the nodes appear.

To set the node order:

1 Right-click any node displayed in the topology map and select **Node Order** to open the Set Node Order dialog box.

The Set Node Order dialog box displays a row for each node that is included in the report.

2 Use the up and down arrows to determine the order of each node included in the report.

3 Click **OK**.

Refining the Report Layout

You can choose to display nodes dependent on other nodes either alongside or beneath the parent node. The latter option can be selected for better visual identification of parent nodes and their dependents.

For example:

➤ The child node (IP) is displayed next to the parent node (Host):

DNS Name	Key	SYS Name	Domain	Net Class	Net Address
10.0.64.110	10.0.64.110 niceDomain		niceDomain	C	10.0.64.0
10.0.64.111	10.0.64.111 niceDomain		niceDomain	C	10.0.64.0
10.0.64.119	10.0.64.119 niceDomain		niceDomain	C	10.0.64.0
10.0.64.120	10.0.64.120 niceDomain		niceDomain	C	10.0.64.0
10.0.64.130	10.0.64.130 niceDomain		niceDomain	C	10.0.64.0
10.0.64.131	00096B2E0BCC	aix-dev	niceDomain	C	10.0.64.0

➤ The child node (IP) is displayed below the parent node (Host):

DNS Name	Key	SYS Name	
10.0.64.110	10.0.64.110 niceDomain		
	Domain	Net Class	Net Address
	niceDomain	C	10.0.64.0
DNS Name	Key	SYS Name	
10.0.64.111	10.0.64.111 niceDomain		
	Domain	Net Class	Net Address
	niceDomain	C	10.0.64.0
DNS Name	Key	SYS Name	
10.0.64.119	10.0.64.119 niceDomain		
	Domain	Net Class	Net Address
	niceDomain	C	10.0.64.0
DNS Name	Key	SYS Name	
10.0.64.120	10.0.64.120 niceDomain		
	Domain	Net Class	Net Address
	niceDomain	C	10.0.64.0

To refine the report layout:

1 On the topology map of the Report Manager, right-click a relationship in the TQL scenario displayed and select one of the following layout options:

➤ Select **Sibling** to display the child node columns next to the parent nodes in the report.

➤ Select **Child** to display the child node columns indented beneath the parent nodes in the report.

2 Repeat step 1 for each additional relationship in the TQL scenario.

Saving the System Report Template

After you have defined the System Report template, you must save it to the database. If you exit the Report Manager without saving, a confirmation message is displayed, requesting that you save your report.

To save the system report template rule:

Display the report to be saved in the topology map, and click **Save** on the toolbar or select a different report or tab in the View Explorer. A message is displayed, asking if you want to save the changes.

The System Report template you have created is saved to the database.

Viewing and Saving the Generated Report

You can display a report in several formats: HTML, table, and Microsoft Excel. In addition, after you generate the report, you can save it to a disk. By saving the report this way, you can view it at a later time, send it to other interested parties and display it independently, without the need to open and use Mercury Application Mapping.

To view and save the report:

1 From the **Configuration Management** menu, select **System Reports**. Double-click the report you want to display or
in the Report Manager, select the report you want to view from the View Explorer, right-click the report, and select the **Generate** option.

2 Select the format in which the report is to be displayed and saved:

➤ **Html** – Enables you to display the report in your browser and save it in HTML format.

➤ **Table** – Enables you to display the report in a table format, and to convert its data into a chart. For details, see "Converting a System Report to a Chart" on page 360.

➤ **Excel** – Enables you to display the report in Excel, and to use Excel functionalities to manipulate the report data. You can save the report in all the formats that are available in Excel.

3 Click **OK** if you are generating the report from **Configuration Management > System Report**.

4 The message: **The requested report is being calculated. You will be notified when it is done.** is displayed. Click **OK**.

5 The following dialog box is displayed.
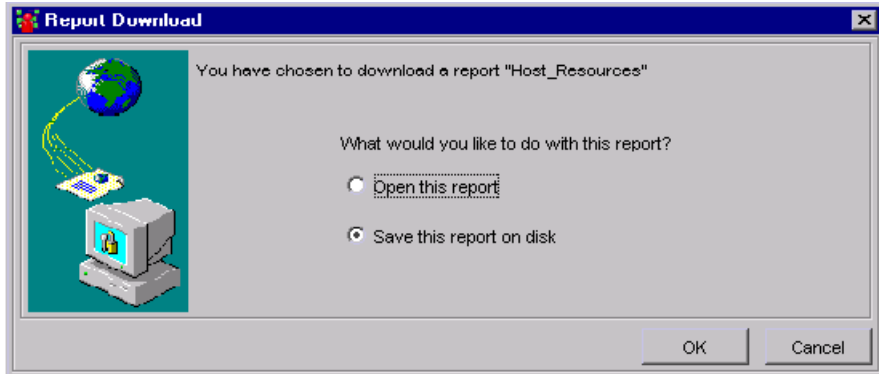
6 Click **Open this report** and click **OK** if you want to open the report or click **Save this report on disk** and click **OK** if you want to save the report. The Save dialog box opens. Enter a unique name for the report and save it in the required location.

➤ If you selected a **Table** format for your report, the report is displayed.

Note: You cannot save a Table format report, but you can view it.

➤ If you selected **Html** or **Excel** format, the Report Download dialog box is displayed.

7 To display the report, select **Open this report** and click **OK**. The report is displayed, as shown in the following HTML System report:

DNS Name	Key	SYS Name	Domain	Net Class	Net Address
10.0.64.110	10.0.64.110 niceDomain		niceDomain	C	10.0.64.0
10.0.64.111	10.0.64.111 niceDomain		niceDomain	C	10.0.64.0
10.0.64.119	10.0.64.119 niceDomain		niceDomain	C	10.0.64.0
10.0.64.120	10.0.64.120 niceDomain		niceDomain	C	10.0.64.0
10.0.64.130	10.0.64.130 niceDomain		niceDomain	C	10.0.64.0
10.0.64.131	00096B2E0BCC	aix-dev	niceDomain	C	10.0.64.0

In the above example, the nodes are displayed in a parent/child relationship with the child nodes indented beneath the parent nodes.

8 To save the report to a disk, select the **Save this report on disk** option button in the Report Download dialog box, and click **OK** to open a Save dialog box.

9 Enter a unique name for the report and save it in the required location.

Importing a System Report

You can import XML files that contain saved system reports to your Report Manager. Use this option to relocate system reports from one workstation to another.

To import a system report:

1 Verify that TQL query attached to a system report appears in the **Attached TQL** list in the New dialog box. If the attached query does not exist in your query list, importing fails.

2 From the Report Manager, open the **Map** menu and click **Import** to open the Import dialog box (for details, see "Importing a TQL Query" on page 91).

3 Locate the report you want to import. Select it and click **Import.** The imported report is added to the system report list in the View Explorer.

Converting a System Report to a Chart

You can convert the data series of one or two report's columns and display it into a chart. By converting the report data to a chart, you can produce additional statistics about the selected data. In addition, the chart's visual representation can help you see and demonstrate comparisons, tendencies and trends in the managed data.

To convert a system report to a chart:

1 Create a system report in a table format and display it. For details, see "Viewing and Saving the Generated Report" on page 357.

2 Right-click the table and select **Chart Manager** to open the Chart Editor dialog box. For details, see "Converting a Map Report into a Chart" on page 238.

3 Click the **Add** button to create a new chart based on the System Report data, and follow the steps described in "Defining a Chart" on page 239.

21

Creating a Sample System Report

This chapter explains how to create a sample system report.

This chapter describes:	On page:
About Creating a Sample System Report	362
Creating the Report TQL Query	362
Creating the Report	365
Defining the Report Nodes	367
Sorting the Report Node Data	370
Setting the Report Node Order	371
Refining and Saving the Report Layout	372
Viewing and Saving the Report	372
Converting the Report to a Chart	373

About Creating a Sample System Report

This section leads you step-by-step through the process of creating a sample System Report. The purpose of this sample report is to display services that are currently malfunctioning. The report also provides details about the services: their location (that is, their hosts), their current state, whether their problems have been acknowledged, the event messages relating to the services' problems, and when the problems occurred.

Creating the Report TQL Query

You first create the report TQL query.

To create the report TQL query:

1 In the TQL Builder, click the **Map** > **New** buttons to open the Create New TQL dialog box.

2 In the **TQL Name** box, type **Services**.

3 From the **Type** list, select **Report**.

4 In the **Description** box, enter **Malfunctioning Services**.

5 Select **Persistent** to define whether the report should always remain in the system memory. You should only use this option for reports that are used frequently.

Note: Reports that are not persistent are recalculated by the system only when they are selected for display. This may take some time, depending on the size of the view.

---

6 Click **OK**. The Services TQL is displayed in the View Explorer.

7 With the Services TQL selected, from the CI Type Model on the right side of
the topology map, click and drag the following Report nodes to the
topology map area:

➤ Network (under Network Resource)

➤ Host

➤ Service (under Host Resource)

The topology map looks something like this:

8 Create a relationship between the network and host Report nodes and
between the host and service Report nodes as follows:

➤ Select the **network** and **host** Report nodes and right-click them.

➤ Select **Add Relationship** to open the Add Relationship dialog box.

➤ Click **Advanced** to open IT World Links hierarchy.

➤ Select **Relationship.**

➤ Select **member** relationship in the IT World Links hierarchy.

➤ Click **OK**.

➤ Repeat this process to define another container relationship between the
service and host Report nodes. For these Report nodes, select the
container relationship from the IT World Links hierarchy.

The Report TQL now displays as follows:
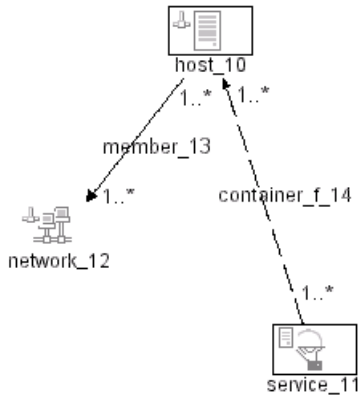
9 The next step is to define the Report TQL to retrieve only services that are indicated as malfunctioning. This is done by defining the attribute condition of the host Report node to retrieve only services that are in a state that is higher than the Normal state. Set this definition as follows:

➤ In the topology map, right-click the **host** Report node and select **TQL Node Definition** to open the TQL Node Definition dialog box.

➤ To define the attribute condition as a state that is higher than Normal, click the **Add** button to open the Condition dialog box.

➤ From the **Attribute Name** list, select **Operation State**.

➤ From the **Operation** list, select **Not equal**.

➤ In the Value comparison area, select **Normal**.

➤ Click **OK**. The TQL Node Definition dialog box now appears as shown in the following figure:
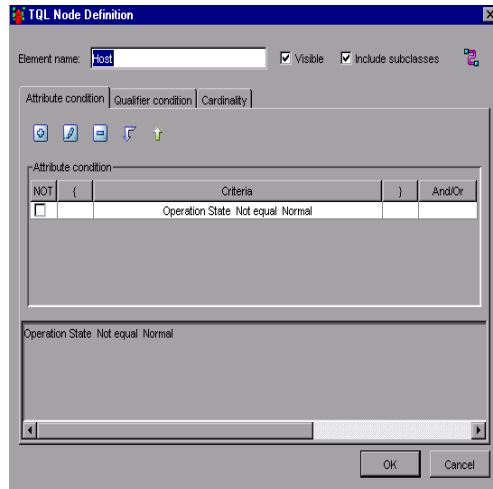
➤ Click **OK**. When Operation State is more than Normal, all services that are higher than Normal are displayed in the report query results.

10 To save the Services TQL, click the **Save** button on the TQL Builder toolbar.

Creating the Report

After defining the query for your report, the next step is to create a report with the Report Manager to display the query results.

To create the report:

1 Click the Report Manager tab.

2 From the Report Manager, click the **Map > New** button on the toolbar to open the Report Definition dialog box.

3 In the **Report Name** box, enter **Services** to define the name of the report in the **System Report** list.

4 In the **Report Description** box, enter **Services with a Normal state**.

5 From the **Attached TQL** list, select **Services**.

6 In the **Report Title** box, enter **Malfunctioning Services**. This title is displayed at the head of the report.

7 In the **Report Sub-Title** box, enter **(Laboratory)**. This sub-title is displayed at the head of the report underneath the title. The dialog box now displays as follows:
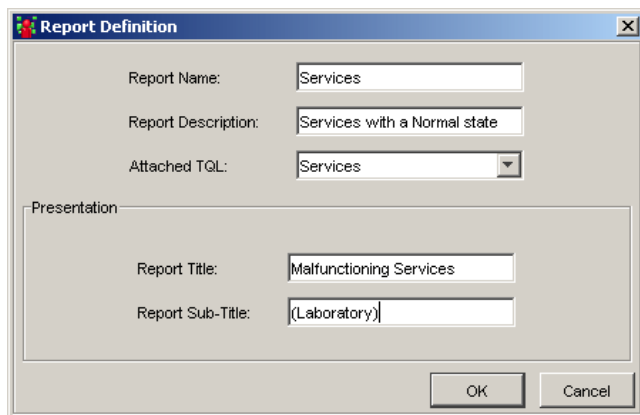
8 Click **OK**. The new report is displayed in the View Explorer and in the topology map on the right side of the window, as shown in the following figure:
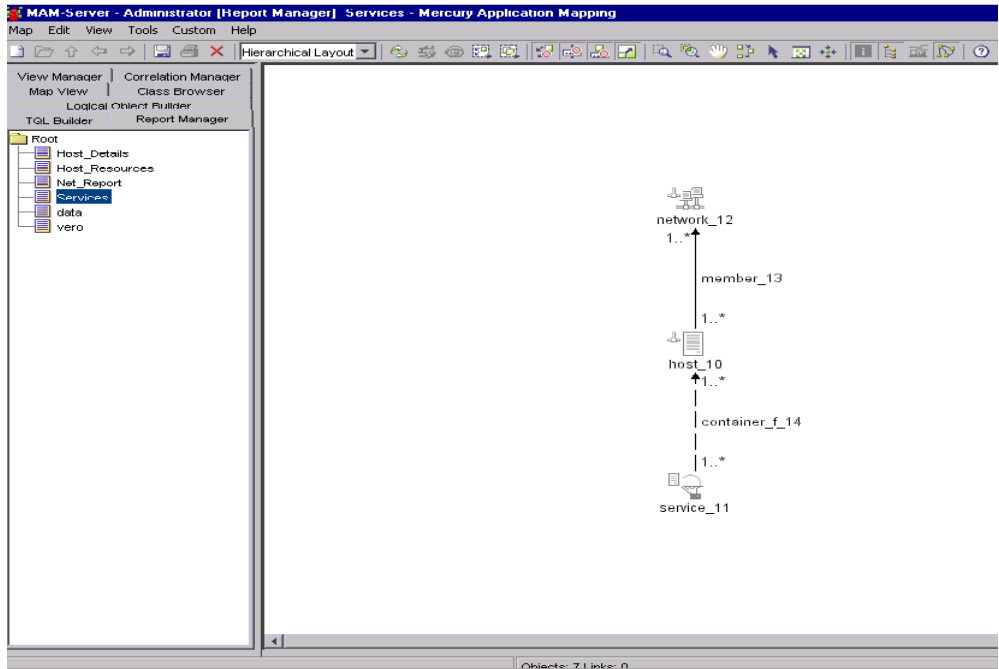
Defining the Report Nodes

After creating the report, the next step is to define the Report nodes that are to be included in the report, that is, the report's columns. This includes determining which information is to be displayed for each Report node (attribute and function), the maximum number of rows to allow for each Report node and the Report node label (that is, the column title).

In this report you have three Report nodes: network, host, and service. For each Report node, you can display the following data in the report's columns:

➤ Network

➤ Host – name

➤ Service – name, functional state and acknowledged state

To define the report nodes:

Beginning with the **network** Report node, select each of the Report nodes in the scenario displayed in the topology map.

1 With the Report node selected, right-click and select **Report Node Definition** to open the Report Node Definition dialog box.

2 Click the **Add** button to display the Column Definition Wizard.

Define the number and content of the selected Report node's columns. You can also enter new titles for the columns. By default, column titles take the same name as the selected attributes or functions, but you can change them.

3 Select **Attribute value columns** and click **Next**.

4 In the **Attribute List** area, select the following attributes for the **network** Report node and click the arrow button:

➤ **Network Domain Name**

➤ **Network Class**

➤ **Network Address**

➤ **State**

5 After you choose an attribute, click > to move it to the **Column List** area.

6 If you wish to change the title of any attribute in the Column List area, click on that attribute. Enter a new title in the Column Title box according to the titles specified in the table below. Click **Apply**. .

Node	Attribute Name	Column Title
Network	network_domain	network_domain
	network_netclass	network_netclass
	network_nettype	network_nettype
	Operation State	Operation State

7 Click **Finish**. You return to the first page of the Column Definition Wizard.

The Report Node Definition dialog box of the network Report node now displays as follows:
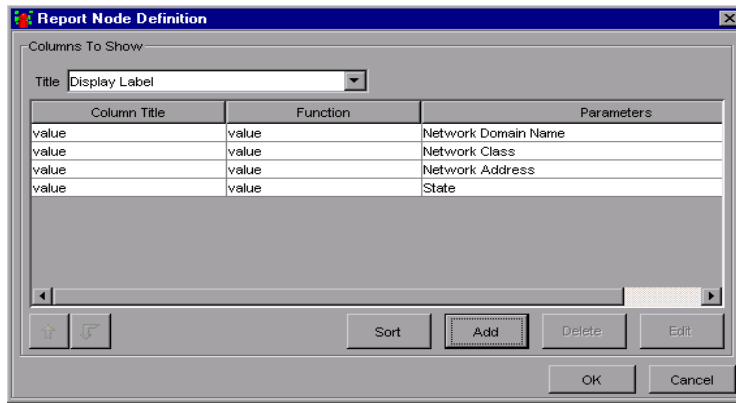
Each row in the table represents a column that will appear in the report.

8 To change the order of the network's columns, use the up and down arrows.

9 Click **OK**. The Report Node Definition dialog box closes.

10 Repeat steps 10 to 5 to add column definitions for the host and service Report nodes. In the Column Definition Wizard, select and enter the following:

➤ For the **host** Report node:

• **Host DNS Name**

• **State**

➤ For the **service** Report node:

• **Name**

• **Service Operating Status**

• **State**

369

Sorting the Report Node Data

You must now determine the order of the information displayed for each Report node, and the method by which the data is sorted in each column.

To sort the Report node data:

Perform the following steps for each Report node displayed in the topology map.

1 Right-click the appropriate Report node and select **Report Node Definition** to open the Report Node Definition dialog box of the **network** Report node.

2 Click the **Sort** button to open the Attribute Columns dialog box.

3 Sort the column information, as follows:

Node	Sort Items by	Then by	Order of Direction
network	network_netclass		Descending
		network_netaddr	Descending
		state	Descending
		network_domain	Descending
host	host_dsname		Descending
		state	Descending
service	Name		Descending
		data_name	Descending
		service_operatingstatus	Descending
		state	Descending

4 Click **OK** to close the Attribute Columns dialog box and to re-display the Report Node Definition dialog box.

5 Click **OK** to close the Report Node Definition dialog box.

Setting the Report Node Order

You must now determine the order in which the Report nodes appear in the report. Since the **network** Report node is the parent Report node, it should appear at the beginning of the report. The data of the **service** Report node is displayed next followed by the events related to the services.

To set the Report node order:

 1 Right-click any Report node displayed in the topology map and select **Node Order** to open the Set Node Order dialog box.

The Set Node Order dialog box displays a row for each Report node that is included in the report.

 2 Use the up and down arrows to move the **network** Report node to the beginning of the list and the **service** Report node to the bottom.

The Set Node Order dialog box now displays as follows:
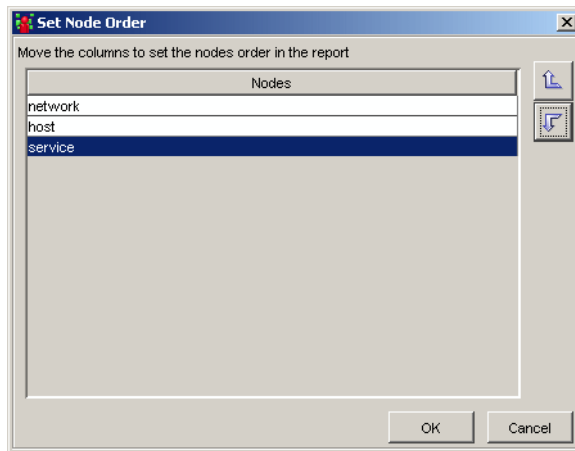
 3 Click **OK**.

Refining and Saving the Report Layout

For better visual identification of the parent hosts and their dependent services, you define the **host** Report node as a child of the **network** Report node. The relationship between the **service** and **host** Report nodes is defined as sibling, to display their data side by side.

To refine and save the report layout:

1 In the topology map, right-click the relationship in between the network and the **host** Report nodes and select **Child** to display the **host** Report node columns indented beneath the **network** Report node in the report.

2 Right-click the relationship between the service and the host and select **Sibling** to display the **service** Report node columns next to the **host** Report node in the report.

3 To save all the report definitions, click the **Save** button.

Viewing and Saving the Report

You can view the report you created and you can save it. You can display the report in three formats—HTML, Table and Excel—and save it in HTML format and in all formats available in Excel.

To view the report:

1 From the **Configuration Management** menu, select **System Reports** to open the System Reports dialog box.

2 From the list of reports, double-click the Services report that you created. The Type dialog box is displayed.

3 Select **HTML** to display and view the report in your browser.

4 Click **OK**. Click **OK** to the confirmation message.

5 Click **OK** to open the Report Download dialog box.

6 To display the report, select the **Open this report** button and click **OK**. The report is displayed, and is similar to the following figure:

Malfunctioning Services **(Laboratory)**				
network_domain	network_netclass	network_nettype	Oper State	
niceDomain	C	iso88023Csmacd	Normal	
	host_dnsname	Oper State	Name	service_operatingstatus	Oper State

			Name	service_operatingstatus	Oper State
			Network Connections	Stopped	Critical
			Removable Storage	Stopped	Critical
			OracleOraHome81HTTPServer	Stopped	Critical
			Background Intelligent Transfer Service	Stopped	Critical
			System Event Notification	Stopped	Critical
			Windows Management Instrumentation	Stopped	Critical
			OracleServiceSKAZAL	Stopped	Critical
	dikla.appilog.com	Normal	SNMP Trap Service	Stopped	Critical
			OracleOraHome81DataGatherer	Stopped	Critical
			Windows Management Instrumentation Driver Extensions	Stopped	Critical
			Hummingbird Jconfig Daemon	Stopped	Critical
			Simple Mail Transport Protocol (SMTP)	Stopped	Critical
			Task Scheduler	Stopped	Critical

7 To save the report to a disk, select the **Save this report on disk** option in the Report Download dialog box, and click **OK** to open the Save dialog box. Enter a unique name for the report and save it in the required location.

Converting the Report to a Chart

You can convert and display the data series of one or two report's columns in a chart. By converting the report data to a chart, you can produce additional statistics about the selected data. In addition, the chart's visual representation can help you see and demonstrate comparisons, tendencies and trends in the managed data. For details, see "Defining a Chart" on page 239.

Part IX

The CI Type Manager

22

Introduction to the CI Type Manager

This chapter introduces the CI Type Manager. The CI Type Manager enables you to view the structure of your managed world.

This chapter describes:	On page:
About the CI Type Manager	378
Working With the CI Type Manager	379
Browsing the CIT Model	380

About the CI Type Manager

The CI Type Manager is displayed by selecting the CI Type Manager tab. The CI Type Manager enables you to view the information in the CI Type model, which contains the definitions of all the configuration item types (CITs) defined in the system and the relationships that define the connections between them. Each CIT has its own attributes, as well as the attributes inherited from its parent CIT. The CI Type Manager represents the structure of Mercury Application Mapping's managed world.

The CI Type Manager visually reflects the data contained in the CMDB, including TQL queries and managed views.
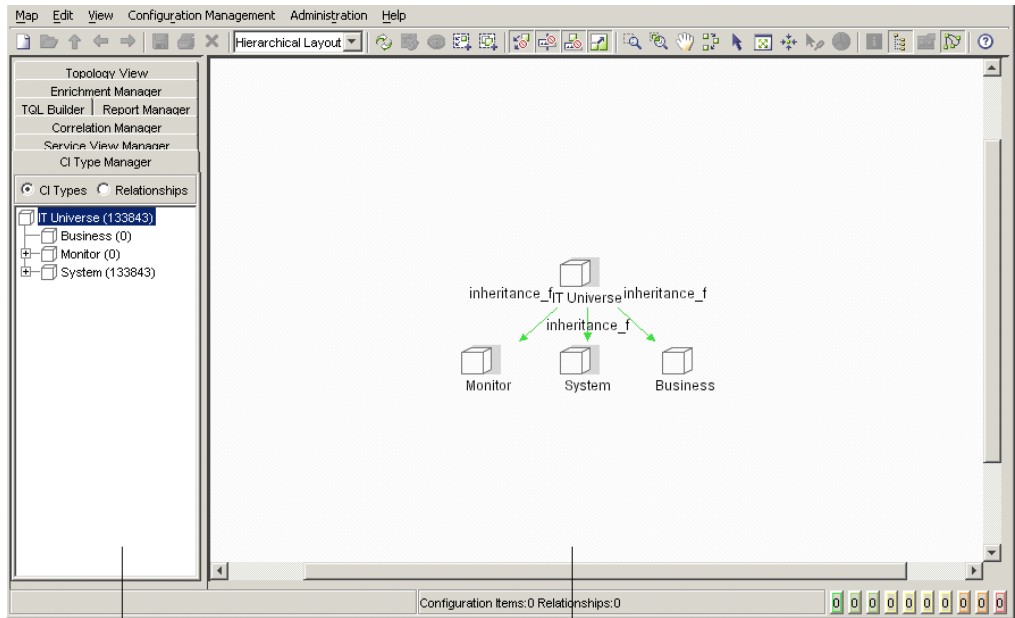
Configuration Item
Types

Topology Map

The CI Type Manager tab is divided as follows:

➤ **Configuration Item Types –** Displays a hierarchical tree structure of the CI Type model containing the inheritance relationships among CITs. All CITs included in the CI Type Model are classified as either a CIT or a relationship. By selecting a CIT in the View Explorer, you can drill down and view the relationships and neighbors of the selected CIT in the topology map. The View Explorer displays the number of instances there are of each CIT in the CMDB.

Note: Each CIT definition type is represented by a unique icon.

➤ **Topology Map –** Displays various layouts of the CIT Model, including the CITs and the relationships connecting them. Each CIT has a unique icon. Its neighbors, as indicated by its connection to the CIT in the View Explorer, all share the same shape surrounding the icon. Shapes are predefined according to logical CIT groupings.

Working With the CI Type Manager

Use the CI Type Manager to perform the following functions:

➤ Create new CITs and edit existing ones.

➤ Expand or collapse the View Explorer to display the required CITs and relationships in the topology map.

➤ View the selected CIT's neighbors, which are connected using inheritance, container and other types of relationships.

➤ Zoom in and out of the topology map to view the selected CIT and its neighbors at different zoom levels.

Browsing the CIT Model

The CI Type Manager enables you to browse the CIT Model to better understand the relationships between specific CITs.

This section includes the following topics:

➤ "CI Type Manager Tooltip" on page 380

➤ "Browsing the CI Type Model" on page 380

➤ "Displaying CIT Instances" on page 382

CI Type Manager Tooltip

When you point to a CIT in the CI Type Manager, the tooltip displays the following information:

➤ **Label –** The name of the CIT to which the item belongs.

➤ **Instances –** The number of instances that currently exist in the CMDB from the selected CIT or its children.

Browsing the CI Type Model

In the View Explorer, you can view inheritance relationships between CITs (such as the **Server Port**, **Client Port**, and **IP Unknown** CITs, which inherit attributes from the **TCP/IP Port** CIT). In the topology map, you can view relationships of all types that connect various CITs. If required, you can print the contents of the topology map.

To browse the CIT Model:

1 From the View Explorer, select the CIT that you want to view in the topology map. The CIT is automatically selected in the topology map.

2 Use either the **Zoom** or **Interactive zoom** button on the toolbar to zoom in on the view area of the selected layer.

3 To return to the default view, click the **Fit To Window** button on the toolbar.

Displaying CIT Instances

You can display all of the instances found in the CMDB for each CIT in a table.

To display all of the instances found for each CIT in a table:

1 In the Topology Map or the View Explorer, right-click the required CIT and select **Show CIT instances** to open the Show all instances dialog box.

The list of CITs are divided into pages. The number at the bottom of the screen indicates which page is currently being displayed. For example, 2/4 means that it is the second out of four pages.

2 To view other pages, use the left and right arrows.

3 To determine the number of node instances that appear on a page, do the following:

➤ Click the **Set bulk size** button to open the Set bulk size dialog box.

➤ Use the up and down arrows or type the number of node instances you want to appear on a page and click **OK**.

4 To see the table updated, click the **Refresh** button.

5 Click **OK** to save the settings you have defined.

23

Managing CITs

This chapter explains how to create and add CITs to the CI Type Model.

Note: The CI Type Manager is not available to Mercury Managed Services customers.

About Managing CITs

If you have IT resources that are not included in the system's built-in CITs, you can create and add CITs to the Configuration Item Type Model. In addition, you can edit existing CITs to adjust them to the structure of your IT infrastructure.

For a list of the existing CITs and their descriptions, see "CI Type Descriptions" on page 499

For a list of the existing relationships and their definitions, see "Relationship Definitions" on page 543.

CI Type Workflow

You create CITs according to the following workflow:

➤ Create CITs. For details, see "Creating CITs" on page 384.

➤ Assign an icon to a CIT. For details, see "Assigning an Icon to a CIT" on page 389.

➤ Create methods for the CIT. For details, see "Creating CIT Methods" on page 391.

➤ Define CIT qualifiers. For details, see "Defining Qualifiers" on page 395.

➤ Define an attribute for the CIT label. For details, see "Defining an Attribute for a CIT Label" on page 396.

Creating CITs

You can create and add CITs to the CI Type Model.

To create a CIT:

1 Select the **CI Type Manager** tab.

2 Click the **New** button on the toolbar or select **New** from the **Map** menu. The Create Configuration Item Type dialog box is displayed.

3 In the **Configuration Item Type Name** box, enter a unique name for the new CIT. Do not use a blank space or an underscore as part of the CIT name. You can use lower and upper case, but you cannot use the same name with different cases for two CITs.

4 In the **Display Name** box, type the name of the CIT as you would like it to appear on the Mercury Application Mapping interface.

5 In the **Base Configuration Item Type** box, select a base CIT for the CIT you are creating. The new CIT inherits the base CIT's attributes.

6 (Optional) In the **Configuration Item Type Description** box, enter a description of the new CIT.

7 Click **Next**. The Attributes page displays all the attributes that the new CIT inherits from its base CIT.

Name	Display Name	Type	Description	Default Valu
FAMILY_ICON	FAMILY_ICON	string		
BODY_ICON	BODY_ICON	string		general
ICON_PROPERTIES	ICON_PROPERTIES	xml		
root_class	CI Type	string	Class name ...	
root_subsystem	Sub System	integer		
root_id	Mam ID	integer		
root_createtime	Create Time	date	When was t...	
root_container	Container	string	Container obj...	
root_system	System	string		
root_updatetime	Update Time	date	When was t...	
root_containername	Container Name	string		
root_uid	root_uid	bytes		
root_lastaccesstime	Last Access Time	date	When was t...	
root_candidatefordele...	Candidate For Deletion...	date	When will thi...	
root_actualdeletetime	Actual Delete Time	date	When will thi...	
root_enableageing	Enable Ageing	boolean	Is ageing en...	false
MENU	MENU	xml		<MENU><ME
data_testcorrstate	Test Corr State	teststates_e...	Test State	Normal
data_teststate	Test State	teststates_e...	Test State	Normal
data_source	Created By	string		
data_name	Name	string		
data_note	Note	string		
data_adminstate	Admin State	adminstates...	Admin State	Managed
data_changestate	Change State	changestate...	Change State	No Change

Create Configuration Item Type-Siebel

Add / Edit / Delete

<<Back / Next>> / Cancel

8 To add attributes, click **Add** to open the Add Attribute page.

9 Enter the following information about the new attribute:

➤ **Attribute Name** – Enter a unique name for the new attribute to identify it in the database.

Note: Do not use a space or an underscore as part of the attribute name.

➤ **Display Name** – Enter a name for the new attribute to identify it on the Mercury Application Mapping interface.

➤ **Description** (Optional) – Enter a description for the new attribute.

➤ In the **Attribute Type** area, select one of the following:

• **Primitive –** Choose from one of the following field types: **boolean**, **bytes**, **date**, **double**, **integer**, **float, long**, **string, xml**.

• **Enumeration/List –** Contains a list of Enumerations/Lists defined in the Enumeration Manager dialog box (for details, see "Enumerations and Lists" on page 67). This option enables you to define an attribute whose value has been predefined.

For example, a location attribute might be defined by a **location** list containing the following values:

• New York

• Boston

• Baltimore

The values that appear in the list in the **Attribute Type** section vary depending on the **Attribute Type** you selected.

➤ **Value Size** – Enter a value for the maximum physical size of the new attribute, if required (depending on the value type you selected).

➤ **Default Value** – Enter a default value for the attribute, if required. This value appears when the new CIT is created and there is no runtime value for the attribute. Choose from the following fields: **Managed**, **Unmanaged**, **Unknown**, **Testing**, **Restricted**, **Disabled**.

➤ In the **Advanced** area:

- (Optional) An index provides rapid access to the attribute values. Select **Index** to accelerate the attribute retrieval performance. This option is recommended for attributes that are used frequently in search conditions. For example, IP address is usually an index attribute of a host.

- (Optional) Select **Set** to define the attribute as a list if the attribute can have multiple values simultaneously. For example, there can be multiple IP addresses for one host.

- (Optional) Select **Required** to define this attribute as a required one, if its value is required for the creation of the CIT.

- (Optional) Select **Visible** to display this attribute in the Attributes page in the topology map.

- (Optional) Select **Editable** to enable future editing of the attribute. Only attributes that are marked as **Editable** (or ones that have values) are displayed in the Attributes page.

➤ In the **Configuration Management only** area:

- (Optional) Select **Change Monitored** to define attributes whose values are being marked as **Change Monitored** in the CI Type Manager. For every change in an attribute defined as **Change Monitored**, you are notified by a Change event in the Topology View. For details on how to track modifications made in the Topology View, see "Viewing and Editing CI Attributes".

- (Optional) Select **Comparable** to enable this attribute to be used for comparing compound CIs. For information on how to compare compound CIs, see "Comparison Capabilities" on page 423.

- (Optional) Select **Asset Data** to display the attribute value in the Asset Report. For information on how to generate an Asset Report, see "Generating an Asset Report" on page 258.

10 Click **OK** to save the changes you have made. The new CIT attribute you created appears in the Attributes page.

> **Note:** If you modify an attribute belonging to its parent, it turns light blue. If you modify an attribute belonging to the CIT itself, it turns dark blue.

 11 If you want to add more attributes, repeat steps 8 to 10.

 12 To define an attribute as a key attribute, click in the left column beside the attribute name. A key icon appears as shown in the following figure:
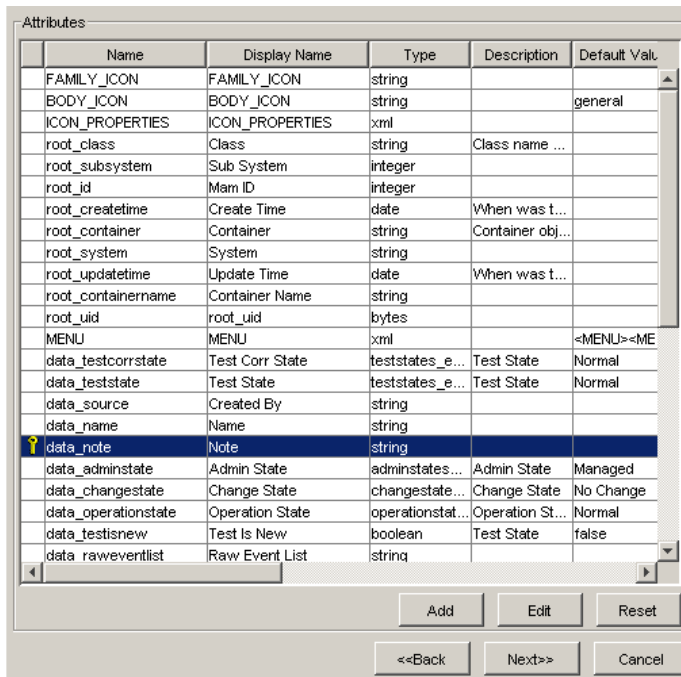
 13 To remove the key attribute definition, click the key icon.

 14 Click **Next** to attach an icon to a new CIT (for details, see "Assigning an Icon to a CIT" below).

Assigning an Icon to a CIT

Mercury Application Mapping allows you to attach an icon to a new CIT. You can also attach different icons to the same CIT when certain conditions apply. For example, you can associate different icons with the same CIT when one of its attribute values changes.

Note: You can only assign different icons to CITs that are classified as CI Types, and not relationships.

To assign one or more icons to a CIT:

 1 In the Attributes page (see step 7 in "Creating CITs" on page 384), click **Next** to open the Icon page.

 2 From the **Configuration Item Type Family** list, select the icon you want to attach to the specified CIT (for details, see "CI Level" on page 41).

 3 From the **Configuration Item Type Main Icon** list, select the group to which the CIT belongs.

 4 Select **Advanced** to assign different icons when the selected CIT attribute gets a different value.

 5 From the **Attribute Name** list, select the attribute to which you want to attach different icons.

 6 Click the **Expand** button to create a new row.

 7 Double-click inside the **Value** column and either select a value from the list that appears, or type the required value.

 8 Double-click inside the **Icon** column and select the icon you want to associate with that value.

 9 To add another value, repeat steps 6 to 8.

In the following example, the computer icon is to be replaced by the NT icon when the attribute's value is **LIKE %NT%,** and by the UNIX icon when the attribute's value is **LIKE %UNIX%**.
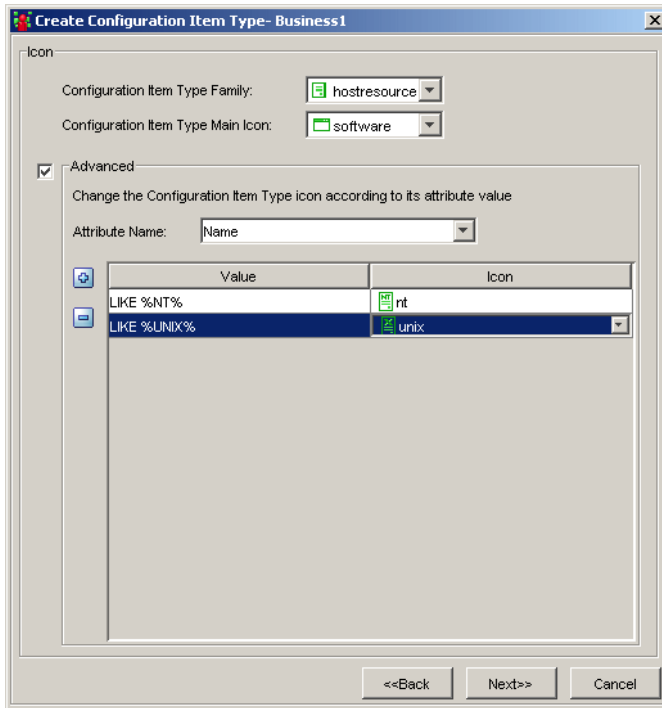
10 To delete a row, select the row you want to delete and click the **Delete** button.

Note: If you change the icon of a CIT that appears in an existing view, the CIT's icon is not updated in the view. For details on how to update the view with the new icon, see "Updating A View" on page 399.

11 Click **Next** to create new CIT methods.

Creating CIT Methods

Mercury Application Mapping enables you to customize the methods (commands) and shortcut menu of a CIT in the CI Type Manager. You can define additional commands, such as ping, run a program, open a URL, and so forth. The newly customized menu is displayed when you right-click the selected CIT in the topology map.

To add a menu item:

1 In the Attribute page (for details on the Attribute page, see "Assigning an Icon to a CIT" on page 389), click **Next** to display the Methods page.
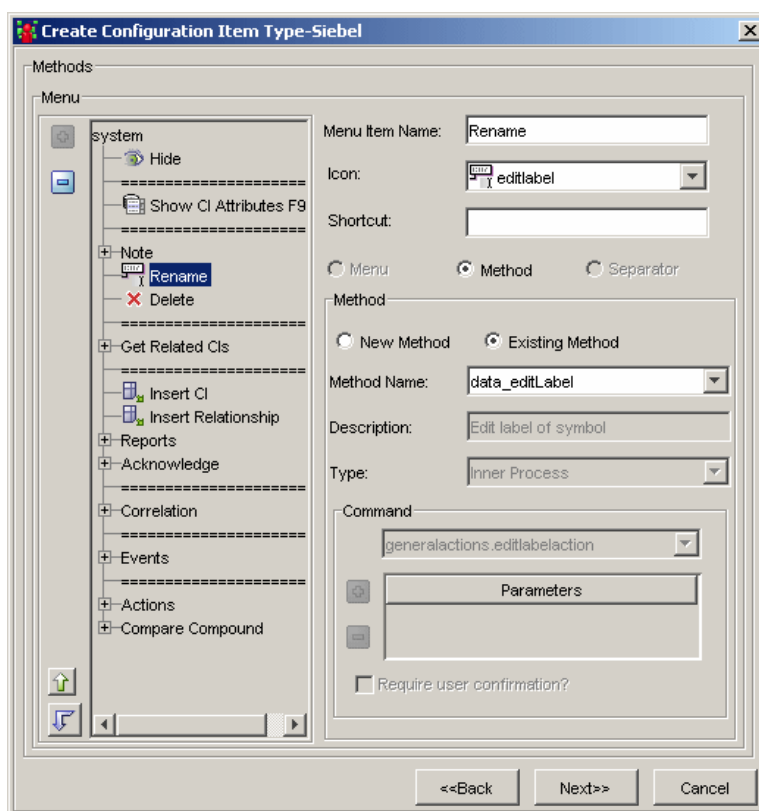
The field on the left displays a hierarchical tree containing the default menu items for the selected CIT.

> **Caution:** It is not recommended to edit the definitions of the default menu items.

 2 Select a menu item from the tree and click the plus button. A new entry appears under the selected item.

3 Select a menu option type, as follows:

➤ **Menu** – Enables you to create a menu option to which you can add a submenu. Continue to step 4.

➤ **Method** – Enables you to add a command to the menu. Skip to step 5.

➤ **Separator** – Places a separator between two menu items. Skip to step 8.

4 To create a menu item to which you can add a submenu, do the following:

➤ Click **Menu**.

➤ In the **Menu Item Name** box, type a name for the menu item as you want it to appear on the menu.

➤ (Optional) From the **Icon** list, select the icon you want to appear next to the menu option.

➤ (Optional) In the **Shortcut** box, press any key combination to create a shortcut for the menu item, for example, CTRL+H.

➤ Click **OK**. The menu option you created is added to the menu in the topology map.

5 To add a command to the menu, do the following:

➤ Click **Method**.

➤ In the **Method** area, select:

• **New Method** to change the command that the menu item runs

• **Existing Method** to choose a command from a list of all the defined methods inherited from the IT World CIT and all its ancestors.

Note: To create a method and not override the existing method in the selected CIT, it is recommended to create a new menu item and place the new method there. To create a new menu item, see step 4.

➤ In the **Method Name** box, type a name for the command.

➤ In the **Description** box, type a description of the command.

➤ From the **Type** list, choose a command type:

• To connect you to a specific location on the World Wide Web, select **URL,** and then in the **URL** box, type the exact Internet address. For example: www.mercury.com.

• To run a program, select **Execute,** and then click the **Find** button to display a standard Browse dialog box and select the program you want to run.

• To display a list of internal Mercury Application Mapping actions select **Inner Process,** and then select the required action.

6 In the **Command** box, type the required command. (This option is available only if you have selected the **URL** or **Execute** command type options.)

If you use a variable in the command field, use the format **command name %1** and define the CIT attributes as described in the next step. The parameter values replace **%1** according to their order in the list. For example, **%1** is replaced by the first parameter in the list, **%2** is replaced by the second parameter in the list, and so forth.

7 To enter CIT attributes in the **Parameters** box, do the following:

➤ Click the plus button to create a new entry.

➤ Type the CIT's attribute.

Note: To delete an existing entry, select it and click the **Delete** button.

8 To place a separator between two menu items, do the following:

➤ Click **Separator**.

9 Click **OK** to place a separator underneath the selected menu option as shown in the illustration below. (Optional) In the **Confirm Message** box,

separator
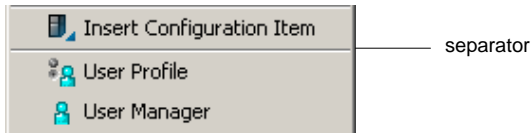
type the message you want to appear when you select the menu item. Select the check box to the left of the **Confirm Message** box to activate this option.

Note: To delete a menu item, select an item and click the minus button.

10 To change the order of the menu items, drag and drop an item into the required position.

Note: If a CIT does not have a method defined specifically for it, the CIT inherits all the menus from its parent CIT or the nearest ancestor that does have a method defined for it.

If you create or modify a menu, the change occurs only in the specific CIT that is edited.

11 Click **Next** to define qualifier conditions for the CIT. For details, see "Defining Qualifiers" on page 395.

Defining Qualifiers

The CI Type Manager allows you to assign qualifiers to a CIT definition. Qualifiers enable you to define added attribute definitions to the CIT. For example, you can use a qualifier to define a CIT as abstract, meaning you cannot create instances from it.

To define qualifier:

1 In the Methods page (for details on the Methods page, see "Creating CIT Methods" on page 391), click **Next** to display the Qualifiers page.

2 In the **Qualifiers** list, select the required qualifiers using the Add buttons to move your selection(s) to the **Configuration Item Type Qualifiers** list. You can make multiple selections by holding down the CTRL key.

> ➤ Use this button to move the selected qualifier(s) to the **Configuration Item Type Qualifiers** list.

>> ➤ Use this button to move all the qualifiers to the **Configuration Item Type Qualifiers** list.

Select the required qualifier(s) according to the following:

➤ **NOTIFY_CHANGES** – For internal use only.

➤ **BLE_LINK_CLASS** – For internal use only.

➤ **ABSTRACT_CLASS** – You cannot create instances of this CIT.

➤ **ITU_HIDDEN_CLASS** – For internal use only.

➤ **CONTAINER** – Relevant for relationships only. Represents containment between two CIs.

➤ **HIDDEN_CLASS** – Does not appear anywhere in the application.

➤ **READ_ONLY_CLASS** – For internal use only.

3 To remove a qualifier, select the qualifier from the **Configuration Item Type Qualifiers** list and click the **Delete** button.

4 Click **Next** to define attributes for a CIT label. For details, see "Defining an Attribute for a CIT Label" on page 396.

Defining an Attribute for a CIT Label

This section describes how to define an attribute for a CIT label. A label is the title that appears under a CI. Label definition can be customized to include different attribute values. For example, if in host the function label is composed of hostname and network, the displayed label is: server1 10.0.65.0.

To define an attribute for the CIT label:

1 In the Qualifiers page (for details, see "Defining Qualifiers" on page 395), click **Next** to display the Default Label page.

2 Select the attribute(s) to appear in the label of the selected item (for details on how to format the text, see "Formatting Text" on page 397 below).

3 Click **Finish**. Mercury Application Mapping creates the new CIT using the information you provided, and displays it on the View Explorer and in the topology map in the Configuration Item Type Model.

Formatting Text

[...] Select the attribute(s) to appear in the formatted text by clicking the button to the left of each attribute to display the selected attribute in the **Format** box. If required, you can use the function buttons to display multiple attributes in the label as follows:

Option	Description
()	Adds parentheses to the formatted text (used in conjunction with the other functions).
&	Places an **AND** operator between two attributes in the formatted text. For example, **network_netaddr&network_domain** displays both the network address and the domain of the node.
|	Places an **OR** operator between two attributes in the formatted text.
RegExp	Adds a regular expression using regular expression syntax to the label definition. For examples of how to use regular expression syntax, see "Using Regular Expressions" on page 398.

Note: To delete an attribute from the **Format** box, highlight it and press the ENTER or DELETE key.

You can define conditions using the following combinations:

➤ AND

➤ OR

➤ AND NOT

➤ OR NOT

Using Regular Expressions

This section gives an example of how to use regular expression syntax.

To enter a regular expression:

1 In the first field, enter the regular expression pattern. This is the structure of the selected attribute.

2 In the second field, enter the group number. This is the part of the regular expression pattern to focus on when creating the label.

For example, enter a regular expression to define the IP address (aa.yy.zz.mm), as follows:

To:	In the first field enter:	In the second field enter:
Create label by zz	(.*[.].*[.])(.*)([.].*)	2
Create label by yy	(.*[.])(.*)([.].*[.].*)	2
Create label by aa	(.*)([.].*[.].*[.].*)	1
Create label by mm	(.*[.].*[.].*[.])(.*)	2

You can also enter a regular expression to create the label by the first or last letter(s) of the selected attribute as the following example shows:

To:	In the first field enter:	In the second field enter:
Create label by the first letter	(.)(.*)	1
Create label by the last letter	(.*)(.)	2
Create label by the first two letters	(..)(.*)	1
Create label by the last two letters	(.*)(..)	2

Updating A View

If you change the icon of a CIT that appears in an existing view, the CIT's icon is not updated in the view. This section explains how to update the view with a new icon.

To update a view with a new icon:

1 Click the Topology View manager.

2 Select the view you want to rebuild from the View Explorer.

3 Click the **Rebuild** button.

Editing Existing CITs

You can edit an existing CIT.

To edit an existing CIT:

1 In the View Explorer, right-click the CIT you want to edit and click **Edit CIT** to open the Edit Configuration Item Type dialog box.

The Edit Configuration Item Type dialog box tabs are similar to those in the Create Configuration Item Type dialog box (as described in "Creating CITs" on page 384). However, only editable fields are enabled in the Edit Configuration Item Type dialog box.

2 Make your changes in the different tabs as follows:

➤ **Information** (see "Creating CITs" on page 384)

➤ **Attributes** (see "Creating CITs" on page 384)

➤ **Icon** (see "Assigning an Icon to a CIT" on page 389)

➤ **Methods** (see "Creating CIT Methods" on page 391)

➤ **Qualifiers** (see "Defining Qualifiers" on page 395)

➤ **Default Label** (see "Defining an Attribute for a CIT Label" on page 396)

Exporting CITs

You can export XML files that contain saved CITs from your CI Type
Manager. Use this option if you want to move CITs from one workstation to
another.

To export a CIT:

1 Right-click the CIT whose XML file you want to export and click **Export** to
open the Save as dialog box.

2 Browse to the location where you want to save the CIT's XML file and click
Save.

Adding Relationships Between CITs

You can add relationships between CITs, either built-in or new ones, which
define their physical or logical connections.

To add a relationship between two CITs:

1 From the View Explorer or the Map pane, select the two CITs you want to
link by holding down CTRL and clicking the CIT names.

2 Right-click one of the CITs and select **Add/Remove Relationship** to open the
Add/Remove Relationship dialog box.

3 Select the check box(es) that specifies the type of relationship with which
you want to link the CITs.

4 Click **OK** to save the changes you have made.

Removing Relationships Between CITs

You can remove relationships between two CITs.

To remove a relationship between two CITs:

1 From the View Explorer or the Map pane, select the two CITs whose relationship you want to delete by holding down CTRL and clicking the CIT names.

2 Right-click one of the CITs and select **Add/Remove Relationship** to open the Add/Remove Relationship dialog box.

3 Clear the check box(es) of the relationships that link the selected CITs.

4 Click **OK** to save the changes you have made.

Part X

Mercury Application Mapping Applications

24

Environment Comparison

This chapter describes how Mercury Application Mapping enables you to take a snapshot of a specific view, save it and then compare it to snapshots taken of that view at different times.

About Environment Comparison

Mercury Application Mapping's environment comparison feature shows the status of a view in one snapshot as opposed to its status taken at a different time. You can compare a current snapshot to a previous snapshot or two snapshots both taken in the past.

This section includes the following topics:

➤ "Taking a Snapshot of a View" on page 406

➤ "Saving Snapshots on Periodic Basis" on page 407

Taking a Snapshot of a View

This section describes how to take a snapshot of a view.

To take a snapshot of a view:

1 In the View Explorer or the topology map of the Topology View manager, right-click the view of which you want to take a snapshot.

2 Select **Snapshot** > **Save** to open the Save Snapshot dialog box.
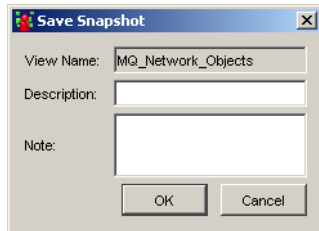
3 The **View Name** box contains the name of the view as it appears in the View Explorer. You cannot edit the view's name.

4 (Optional) In the **Description** box, enter a description of the snapshot.

5 (Optional) In the **Note** box, add a note you want to add about the snapshot.

6 Click **OK** to save the snapshot you have taken in the database. If the snapshot was saved successfully, the following message appears:

Snapshot for view was saved successfully

Saving Snapshots on Periodic Basis

You can schedule Mercury Application Mapping to take snapshots of a certain view at specified times by defining a task in the Scheduled Actions Manager.

To save snapshots on a periodic basis:

1 Select **Administration** > **Scheduler** to open the Scheduled Actions Manager.

2 Click **Add** to define the new task.

3 Click **Add** to open the Actions dialog box.

4 Select **Save a View Snapshot**.

5 Click **Next** to select the view in which you want to save the snapshots.

6 Click **Next** to display a list of parameters to define.

7 Click **Finish**. The action you have defined appears in the Scheduled Action Wizard.
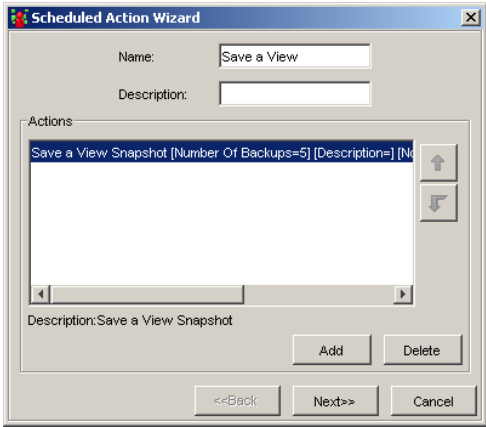
For information on how to schedule the task, see "Task Scheduling" in the *Mercury Application Mapping Administration Guide*.

Comparing Snapshots Taken at Different Times

Mercury Application Mapping allows you to compare two snapshots of a specific view taken at different times. This feature enables you to visualize the differences between the views by comparing the statuses of the view at the times the snapshots were taken.

To compare different snapshots of the same view:

1 Right-click the view of which you have taken snapshots.

2 Select **Snapshot** > **View snapshots** to open the Snapshots dialog box.
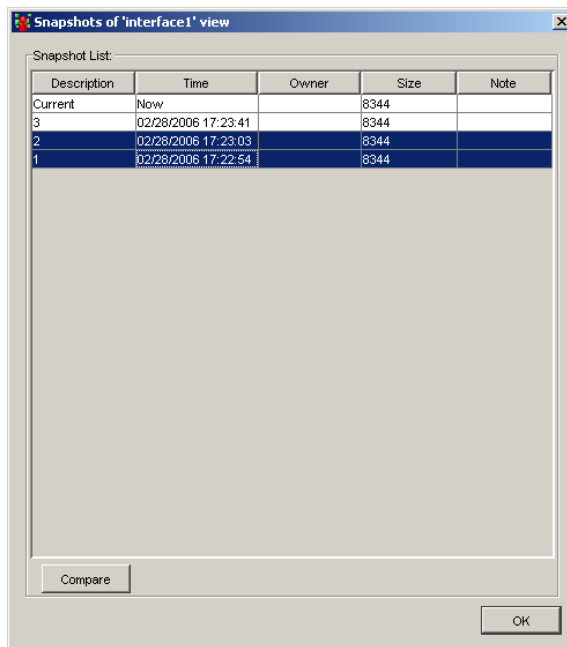
This dialog box contains a list of all the different snapshots taken of the selected view.

The following table describes the fields in the dialog box.

Field	Description
Description	The description you wrote for the snapshot in the Save Snapshot dialog box (see "Taking a Snapshot of a View" on page 406).
Time	The time at which the snapshot was saved.
Owner	The name of the user who took the snapshot.
Size	The number of CIs that were in the view when the snapshot was taken.
Note	The note you wrote for the view in the Save Snapshot dialog box (see "Taking a Snapshot of a View" on page 406).

Note: The first row always represents the status of the view in its current state.

3 Select the two snapshots you want to compare.

4 Click the **Compare** button to open the Compare snapshots dialog box.
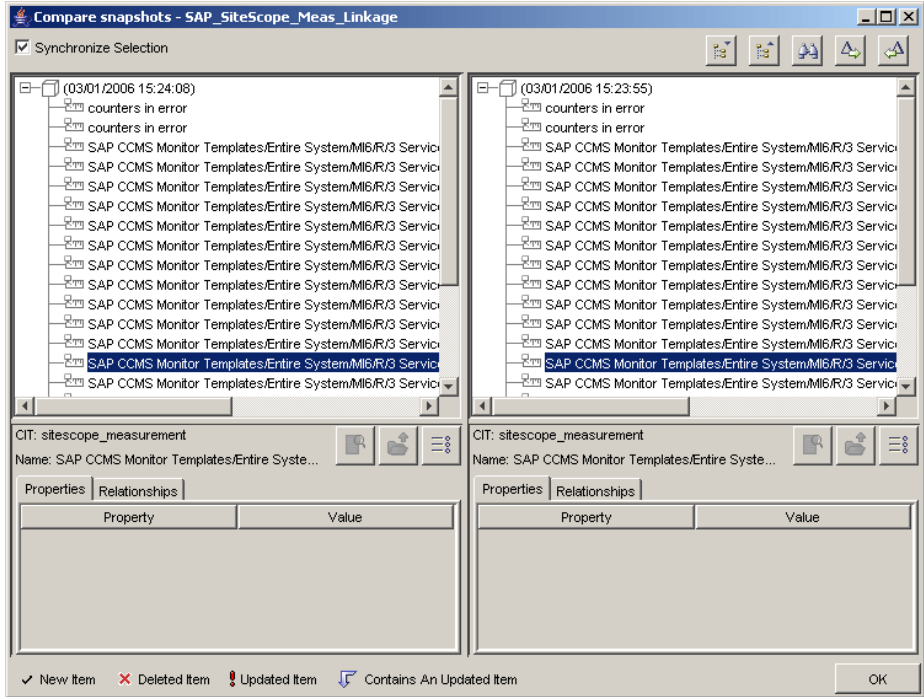
Note: For details about the Compare snapshots dialog box, see "Understanding the Compare Snapshots Dialog Box" on page 410.

Understanding the Compare Snapshots Dialog Box

The Compare Snapshots dialog box is divided into two sections. Each side displays an hierarchical tree structure of all the CIs in the topology map for that snapshot version. The left side always represents the newer of the two versions. When you select a CI on one side, the corresponding CI on the other side is automatically selected (provided the **Synchronize Selection** check box is selected).

The following information is displayed under the tree when you select a CI:

➤ **CIT** – The name of the CIT to which the CI belongs as defined in the CIT Model.

➤ **Name** – The name of the CI as defined in the CI's label.

The rest of this section contains the following topics:

➤ "Compare Snapshots Dialog Box Tabs" on page 412

➤ "Expanding and Collapsing the Entire Tree" on page 412

➤ "Finding Items Marked as Changed" on page 413

➤ "Finding a CI by Its Label" on page 413

➤ "Change Indicators" on page 414

➤ "Viewing and Comparing a Configuration File" on page 417

➤ "Displaying a CI's Attributes" on page 420

➤ "Enabling Synchronized Selection" on page 421

Compare Snapshots Dialog Box Tabs

The Compare snapshots dialog box contains the following two tabs:

Tab	Description
Properties	Displays only the attributes of the CI whose values have been defined as either **Change Monitored**, **Asset Data**, or **Comparable** in the Add Attribute dialog box in the CI Type Manager. For details, see "Creating CITs" on page 384.
	The **Properties** tab has the following two fields:
	➤ **Property** – Displays the attribute name for the selected CI which was marked as **Change Monitored** in the CI Type Manager.
	➤ **Value** – Displays the value for the currently selected CI.
Relationships	Displays the relationships with which the different CIs are connected within the view.
	The **Relationships** tab has the following two fields:
	➤ **Property** – Displays the relationship types. For example, **container_f** or **contained**.
	➤ **Value** – Displays the direction indicator (an arrow either pointing to the left or the right) and the target/source CI to which the arrow is pointing.

Expanding and Collapsing the Entire Tree

This section describes how to expand and collapse the entire tree.

To expand the entire tree:

Select the view at the top of the tree and click the **Expand All** button.

To collapse the entire tree:

Select the view at the top of the tree and click the **Collapse All** button.

Finding Items Marked as Changed

This section describes how to find the items that have been marked as changed.

To find the next occurrence of an item marked as changed:

Click the **Next Difference** button.

To find the previous occurrence of an item marked as changed:

Click the **Previous Difference** button.

Finding a CI by Its Label

Mercury Application Mapping enables you to search for a CI within the selected view whose label contains the searched text. You can specify a whole word or part of it.

To find the CI by its label:

 1 Click the **Find** button to open the Find CI by label dialog box.
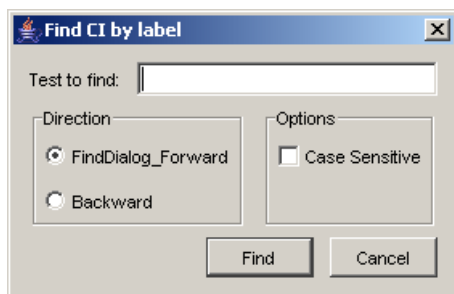
 2 In the **Text to find** box, type the text you want to find.

 3 In the **Direction** section, specify whether you want to do a forward or backward search.

 4 In the **Options** section, select the **Case Sensitive** box if you want the matches to be case sensitive.

 5 Click **Find**. If Mercury Application Mapping finds a match, it selects the CIs that appear in the topology map of both snapshot versions.

Change Indicators

This section describes the icons that indicate the kind of change that has occurred in the compared views. The indicator key appears at the bottom of the Compare snapshots dialog box.

New Item Icon

The **New Item** icon indicates that a new CI has been added to the view. The example below shows that the CI **aPPaPP** has been added to the view in the snapshot at the left.
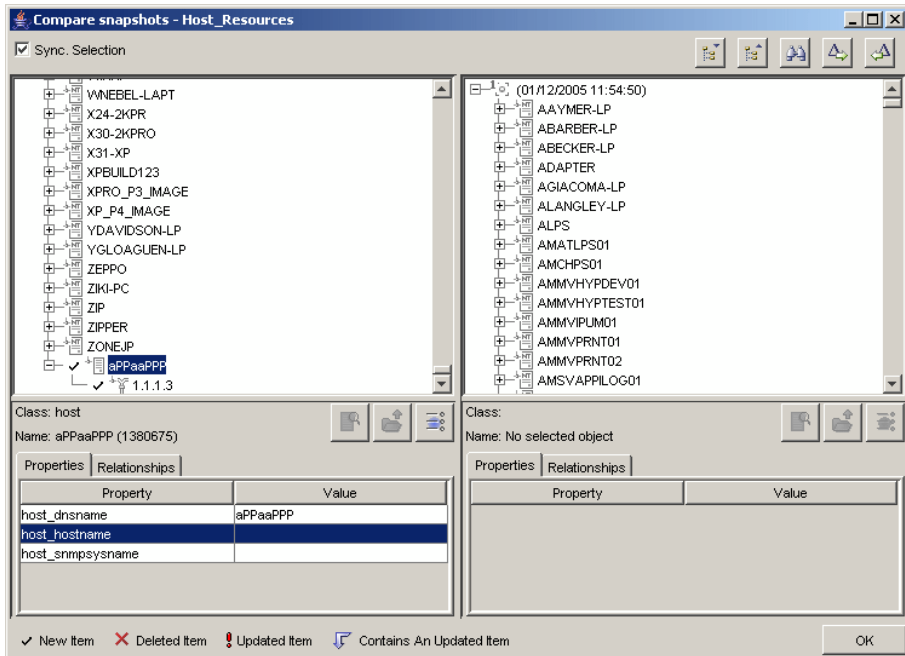
Deleted Item Icon

The **Deleted Item** icon indicates that a CI has been deleted from the view. For example, the figure below shows that the CI **DR2** has been deleted from the view in the snapshot at the right.
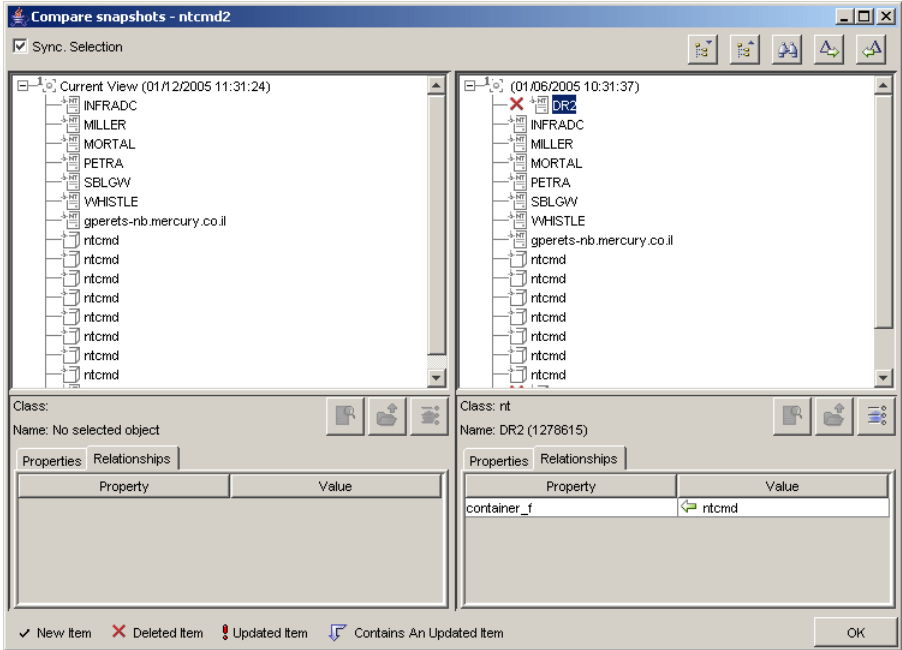
Updated Item Icon

The **Updated Item** icon indicates that a CI whose attribute value has been defined as **Change Monitored**, has changed For example, the figure below shows that the CI called **SQL_APP_Dis_db2** has the **Updated Item** icon because it received the new value **Testing** (see the **Properties** tab underneath the tree structure on the left side).
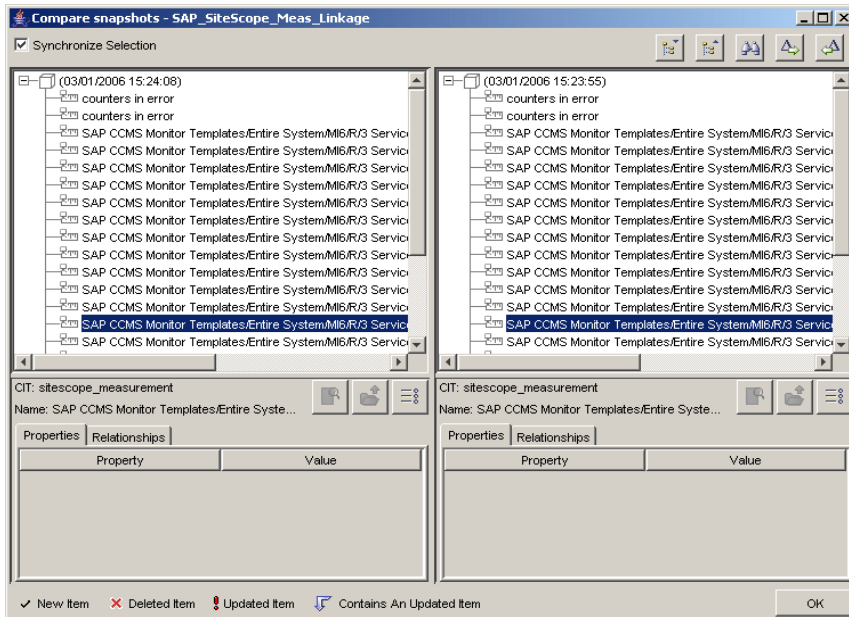
Contains an Updated Item Icon

The **Contains an Updated Item** icon indicates that either a new child CI was created for a CI or an existing child CI was edited or deleted. For example, the **DB2** CI in the figure above has a **Contains an Updated Item** icon due to the changes made to its child CIs.

Viewing and Comparing a Configuration File

A configuration file is a file that has the content of a configuration file, a parameter file, or an **.ini** file. The Compare Snapshots dialog box enables you to view a configuration file, and also compare what the configuration file looked like at different times.

To view a configuration file:

1 In the Compare Snapshots dialog box, select a CI of the CIT **configfile** from the tree.
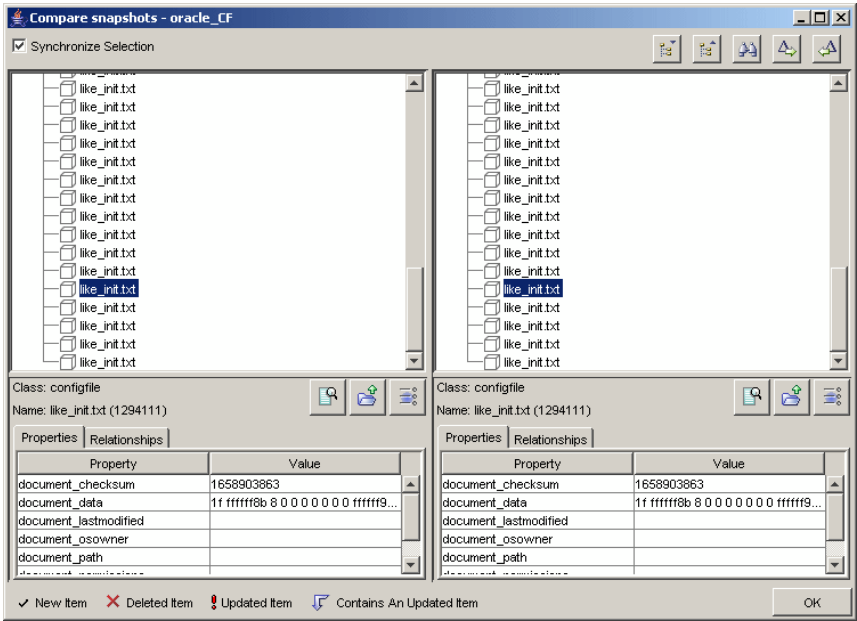
   2 Click the **Open** button to open the configuration file.

Note: This button is only enabled if you have selected a CI of the CIT **configfile** from the tree.

To compare what the configuration file looked like at different times:

1 In the Compare Snapshots dialog box, select a CI of the CIT **configfile** from the tree.

2 Click the **Show Difference** button to display what the configuration file looked like at the time each snapshot was taken.

Note: This button is only enabled if you have selected a CI of the CIT **configfile** from the tree.

The differences in the configuration file are marked in blue.

Displaying a CI's Attributes

You can display what the CI's attributes were at the time the snapshot was taken.

To display a CI's attributes:

1 In the Compare Snapshots dialog box, select a CI from one of the trees.

2 Select the **Show CI Attributes** button to open the Attributes dialog box.
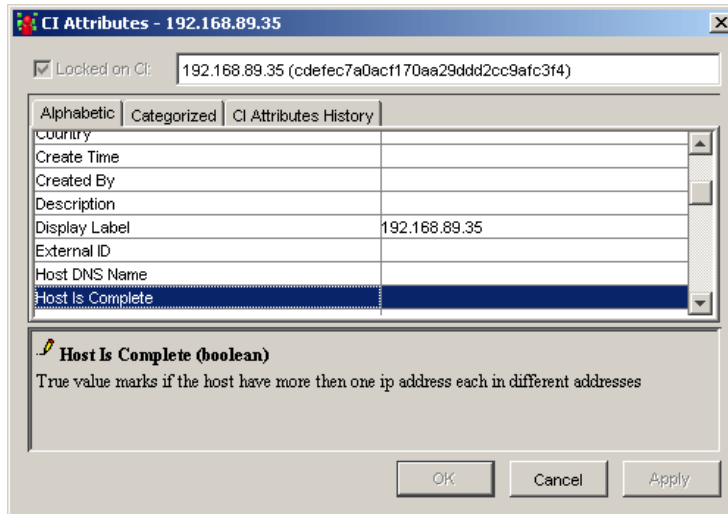
Note: This dialog box is for viewing only.

The Attributes dialog box lists the CITs contained in this view and their attribute values.

➤ The CITs in the **Alphabetic** tab are listed in alphabetical order.

➤ The CITs in the **Categorized** tab are grouped according to the categories defined in the CI Type Manager. Click the **Expand** button to view all the CITs contained in this view

➤ The **CI Attributes History** tab displays a list of attributes in which changes have occurred.

Note: The CI Attributes History tab only displays the attributes whose values have been defined as **Change Monitored** in the Add Attributes dialog box in the CI Type Manager. For details, see "Creating CIT Methods" on page 391.

The CI Attributes History tab has the following fields:

Field	Description
Attribute	The name of the attribute.
Change Date	The date when the change occurred.
Changer	Indicates the cause of the change. For example, the field can contain the name of a user or a discovery pattern.
New Value	The new value of the attribute.

You can reduce the number of attributes that appear in the CIs Attributes History tab by selecting an attribute from the **Filter by** list.

Note: The CI Attributes History tab is only updated every ten minutes.

Enabling Synchronized Selection

Select the **Synchronize Selection** check box at the top-left corner of the Compare snapshots dialog box to enable you to select a CI on one side and have its corresponding CI automatically selected on the other side.

By default, **Synchronize Selection** is selected. Clearing this check box allows you to compare two different CIs by selecting one from each of the two trees.

25

Comparison Capabilities

This chapter describes Mercury Application Mapping's comparison capabilities.

This chapter describes:	On page:
Gold Master Comparison	423
CI Comparison	431

Gold Master Comparison

Mercury Application Mapping version enables you to generate a report that compares the configuration of a Gold Master CI to other CIs of the same CIT.

You choose the CI you want to use as the Gold Master CI. Then you choose the CIs whose configuration you want to compare to the Gold Master template. You can choose the CIs from existing views or from the CMDB.

In the report that is generated, Mercury Application Mapping specifies which of the selected CIs are configured differently.

Note: The CIs in this report will only appear if their attributes are marked as **Comparable** in the Edit Attribute dialog box (see "Creating CITs" on page 384.)

This section includes the following topics:

➤ "Comparing the Configuration of CIs to a Gold Master Template" on page 424

➤ "Deleting a Gold Master Report" on page 428

➤ "Editing a Gold Master Report" on page 429

➤ "Displaying CI Attributes" on page 429

➤ "Searching for a CI" on page 429

➤ "Displaying Related CIs" on page 431

Comparing the Configuration of CIs to a Gold Master Template

This section describes how to compare the configuration of specific CIs to a Gold Master template.

To compare the configuration of CIs to a Gold Master template:

1 In Topology View, select **Configuration Management** > **Gold Master Reports** to open the Gold Master Report Manager window.
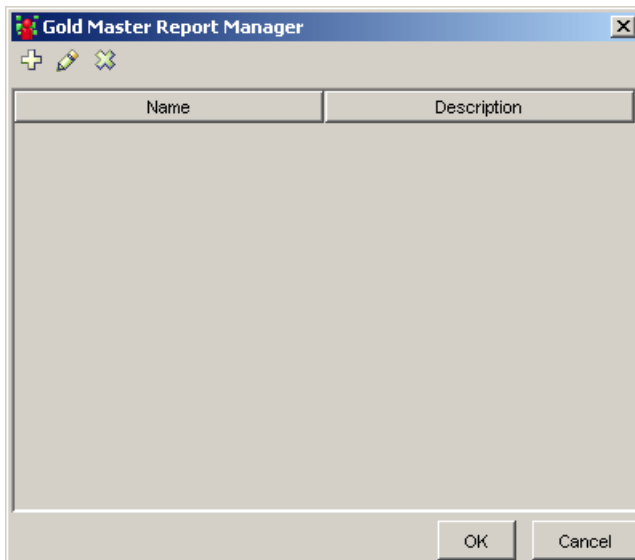
2 Click the **Add** button to open the Gold Master Report window.
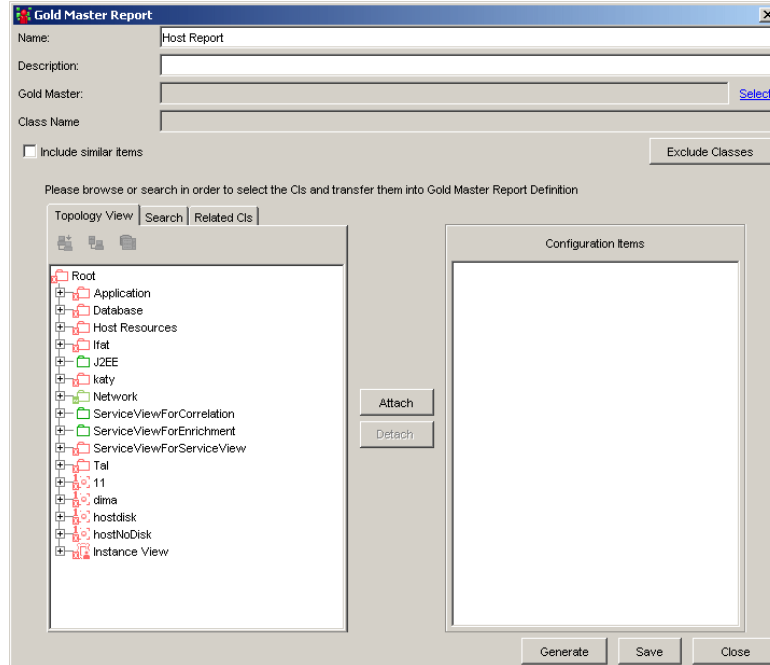
3 In the **Name** box, type a unique name for the Gold Master report.

4 (Optional) In the **Description** box, enter a description of the Gold Master report.

5 To select the CI you want to use as a Gold Master CI, click **Select** at the right end of the **Gold Master** box. The Topology View window opens.

➤ In the **Topology View** tab, select the required CI(s) at different levels of the view hierarchy. You can select multiple CIs by holding down the CTRL button and clicking the required CIs. To view and edit a CI's attributes, see "Displaying CI Attributes" on page 429.

➤ In the **Search** tab, search for CIs in existing views or in the entire CMDB. For details, see "Searching for a CI" on page 429.

➤ In the **Related CIs** tab, select a CI that is related to the selected CI in the **Topology View** tab. For details, see "Displaying Related CIs" on page 431.

6 Select the required CI. The selected CI appears in the **Gold Master** box and the CIT of the Gold Master CI appears in the **CIT Name** box.

7 Select the **Include similar items** check box if you want to include the compared CIs whose configuration is identical to that of the Gold Master CI.

8 To exclude a specific CIT from the Gold Master report, click **Exclude CITs** to open the Exclude CIT window.

9 From the Exclude CIT window, select the required CIT and click **OK**.

10 In the left pane, select the CI(s) that you want to compare to the Gold Master CI. You can select the CIs in the following ways:

➤ In the **Topology View** tab, select the required CI(s) at different levels of the view hierarchy. You can select multiple CIs by holding down the CTRL button and clicking the required CIs. To view and edit a CI's attributes, see "Displaying CI Attributes" on page 429.

➤ In the **Search** tab, search for CIs in existing views or in the entire CMDB. For details, see "Searching for a CI" on page 429.

➤ In the **Related CIs** tab, select a CI that is related to the selected CI in the **Topology View** tab. For details, see "Displaying Related CIs" on page 431.

11 Use the **Attach** and **Detach** buttons to transfer to and remove CIs from the **CI Items** list. You can also select multiple CIs. The CIs that appear in the **CI Items** list are the CIs that are compared to the Gold Master CI.

12 Click **Save** in the Gold Master Report window to save the Gold Master report definition. The Gold Master report definition appears in the Gold Master Report Manager window.

13 To generate a report, click **Generate.** The Type dialog box opens. Define the format in which you want to receive the report. You can receive the report in the following formats:

➤ **HTML**

➤ **Excel**

14 Click **OK**. The following message appears:
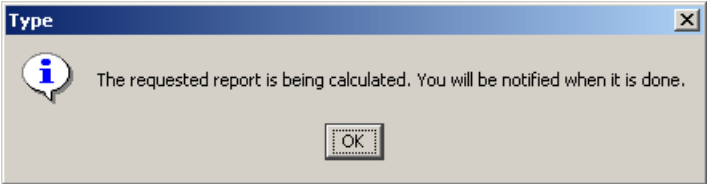
15 Click **OK** to create the report.

Gold Master Object - heroism.mercury.co.il				
Report Time - Jan 19, 2005 1:57:04 PM				
PROWLER				
Component	Parent Object	Count/Gold	Missing	Additional
nt(PROWLER)		1/1		host_os:Windows 2000
				host_vendor:Microsoft
interface	PROWLER	2/1		00112F161C9F()
interfaceindex	PROWLER	1/0		65539_Intel(R) PRO/100 VE Network Connection()
ipserver	PROWLER	8/5		rmi_1099 (Name:rmi,ipport_number:1099) weblogicSSL_7002 (Name:weblogicSSL,ipport_number:7002) weblogic_7001 (Name:weblogic,ipport_number:7001)
j2eeserver	PROWLER	1/0		appserver(j2eeserver_version:WebLogic Server 8.1 SP2 Fri Dec 5 15:01:51 PST 2003 316284 WebLogic XMLX Module 8.1 SP2 Fri Dec 5 15:01:51 PST 2003 316284)
oracle	PROWLER	2/1		globalid()
osuser	PROWLER	6/0		SUPPORT_388945a0()
				HelpAssistant()
				Administrator()
				winroot()
				ASPNET()
				Guest()
snmp	PROWLER	1/0		snmp()
dbaobjects	globalid	1195/1191		STATSPACK(Name:STATSPACK)
				STATSPACK(Name:STATSPACK) DBMS_SHARED_POOL (Name:DBMS_SHARED_POOL) DBMS_SHARED_POOL (Name:DBMS_SHARED_POOL)
dbdatafile	globalid	14/14	D:\ORACLE\ORADATA\SKAZAL\XDB01.DBF (Name:D:\ORACLE\ORADATA\SKAZAL\XDB01.DBF)	D:\ORACLE\ORADATA\SKAZAL9 \XDB01.DBF (Name:D:\ORACLE\ORADATA\SKAZAL9 \XDB01.DBF) D:\ORACLE\ORADATA\SKAZAL9

The Gold Master report has the following fields:

Field	Description
Gold Master CI	The CI that is used as the template to which to compare the other CIs.
Report Time	The time the report was created.
Component	The CIT to which the CI belongs.
Parent CI	The CI under which the compared CI is located.
Count/Gold	Number of CIs in the compared CIs as opposed to the number of CIs in the Gold Master CI.
Missing	Lists the CIs that exist in the Gold Master CI but not in the compared CI.
Additional	Lists the CIs that exist in the compared CI but not in the Gold Master CI.

Deleting a Gold Master Report

This section describes how to delete a Gold Master report from the Gold Master Report Manager window.

To delete a Gold Master report:

1 Select **Configuration Management** > **Gold Master Reports** to open the Gold Master Report Manager window.

2 Select the Gold Master report you want to delete and click the **Delete** button.

3 Click **OK** to save the changes you have made.

Editing a Gold Master Report

This section describes how to edit a Gold Master report from the Gold Master Report Manager window.

To edit a Gold Master report:

1 Select **Configuration Management** > **Gold Master Reports** to open the Gold Master Report Manager window.

2 Select the Gold Master report you want to edit and click the **Edit** button to open the Gold Master Report window.

3 Edit the report as required. For details, see "Comparing the Configuration of CIs to a Gold Master Template" on page 424.

4 Click **OK** to save the changes you have made.

Displaying CI Attributes

This section describes how to display and edit a CI's attributes.

To display and edit a CI's attribute:

1 In the Gold Master Report window, select the **Topology View** tab.

2 Select the CI whose attributes you want to view.

3 Click the **Show CI's attributes** button to open the CI Attributes dialog box. The CI Attributes dialog box enables you to view and edit the attributes defined for the selected CI. For details, see "Viewing and Editing CI Attributes" on page 201.

Searching for a CI

Mercury Application Mapping search capabilities enable you to search for CIs to compare to the Gold Master CI either in specific views or in the entire CMDB. Theses capabilities also provide different search criteria, through which you can search for CIs according to their CIT, label, and/or attributes.

To search for CIs:

1 In the Gold Master Report window, select the **Search** tab.

2 From the **CIT** list, select the CIT of the CI(s) you want to find.

3 (Optional) Select **Derived** to find both the CI(s) of the selected CIT and any CIs derived from the selected CIT (inheritance CIs).

4 (Optional) In the **Options** area, if you want to find a specific CI(s) whose label you know, enter the CI's label in the **Label** box. Once you have entered a label, you can use the following search options:

➤ **Case Sensitive** – Distinguishes between uppercase and lowercase characters. When you select **Case Sensitive**, Mercury Application Mapping finds only those instances in which the capitalization matches the text you typed in the **Label** box.

➤ **Find whole words only** – Searches for occurrences that are whole words and not part of a larger word.

➤ **Starts With** – Searches for occurrences that begin with the entered text.

➤ **Ends With** – Searches for occurrences that end with the entered text.

Note: The options **Starts With** and **Ends With** are only available if you have not selected **Find whole words only**.

---

5 In the **Search In** area, select the place in which you want to conduct the search:

➤ **My Views** – Searches only in the views that belong to your user profile.

➤ **Database** – Searches in the entire CMDB.

6 Click **Search**. The search results are displayed in the **Search** tab.

7 To do another search, click **Back to Search**.

Displaying Related CIs

Mercury Application Mapping enables you to display the interdependencies of any selected CI (for details, see "Displaying Interdependent CIs" on page 205). You can select CIs to compare to the Gold Master CI from a list of related CIs.

You can display:

➤ All CIs connected to a selected CI that are part of a specific view.

➤ All CIs connected to the selected CI in the entire CMDB.

You select a CI in the **Topology View** tab, and the CIs related to the selected CI appears in the **Related CIs** tab.

To display the related CIs of a selected CI:

1 In the Gold Master Report window, select the **Topology View** tab.

2 Select the CI whose related CIs you want to display.

3 To display related CIs, do the following:

➤ Click the **Get Related CIs from View** button to display related CIs in the selected view.

➤ Click the **Get Related CIs from Database** button to display related CIs in the CMDB.

The related CIs are displayed in the **Related CIs** tab.

CI Comparison

This section contains the following topics:

➤ "Comparing Compound CIs" on page 432

➤ "Comparing CIs by View" on page 435

➤ "Understanding the Compare CIs Dialog Box" on page 436

➤ "Understanding the Properties Tab" on page 436

➤ "Expanding/Collapsing the Tree" on page 436

➤ "Finding Items Marked as Changed" on page 437

➤ "Finding a CI by Its Label" on page 437

➤ "Understanding Change Indicators" on page 438

➤ "Viewing and Comparing Configuration Files" on page 440

➤ "Displaying a CI's Attributes" on page 443

➤ "Enabling Synchronized Selection" on page 445

Comparing Compound CIs

A compound CI is a CI that is linked to another CI by a **Container link**. Comparing two compound CIs enables you to see the difference between the hierarchies of the two CIs. For example, Host A may contain only one disk whereas Host B may contain three disks. In addition, you can also see differences in the attributes defined as **Comparable** in the CI Type Manager (see "Creating CITs" on page 384). For example, there might be a difference in the operating system or disk capacity.

Note: You can only compare two CIs from the same CIT.

When you compare two compound CIs, you can either first select the CIs from the Map pane and then open the Compare compound CIs window, or, alternatively, open the Compare compound CIs window first and then select the two CIs you want to compare.

Mercury Application Mapping also allows you to display the child CIs of the compared CIs that appear in the selected view. For details, see "Comparing CIs by View" on page 435.

To compare compound CIs:

1 To compare compound CIs, select **Configuration Management** > **Compare Compound** to open the Compare compound CIs from DB dialog box.
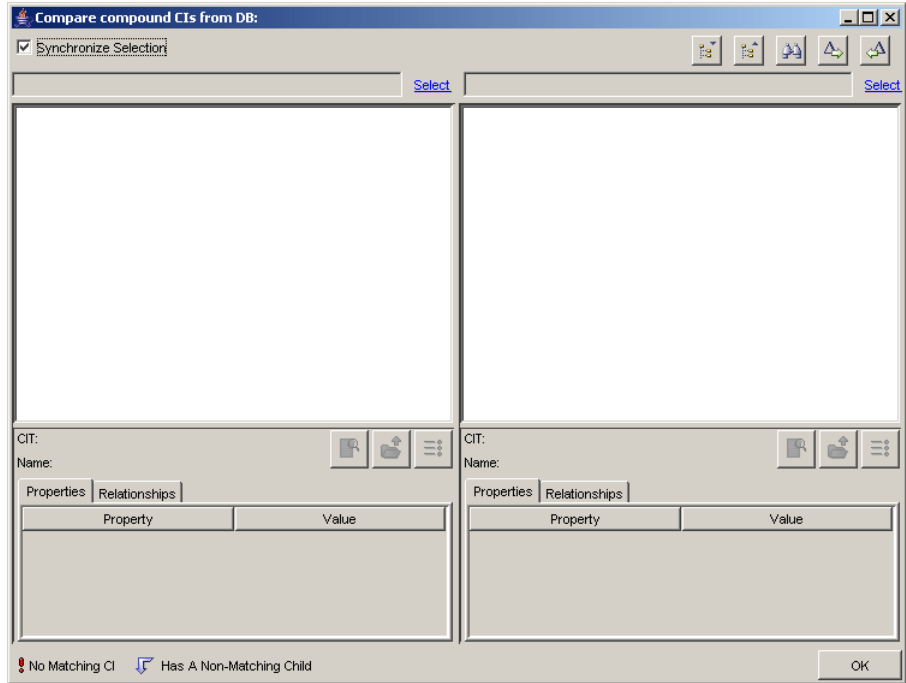
Select

2 Click **Select** above the top-right corner of each pane to select the CIs you want to compare. The following is displayed.
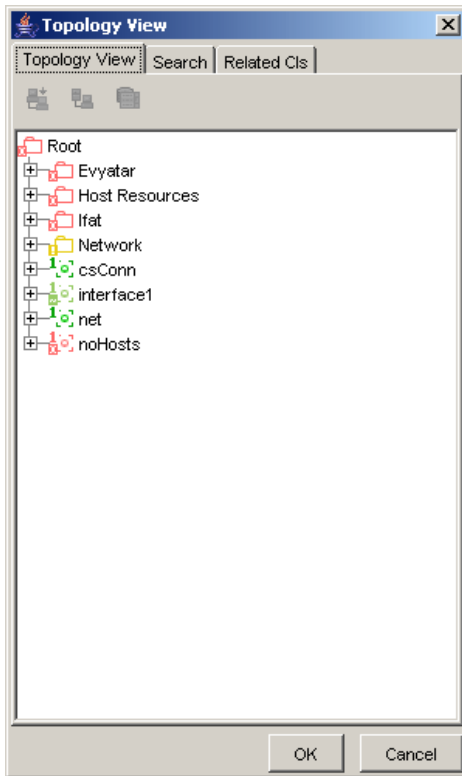
You can select the CIs in the following ways:

➤ In the **Topology View** tab, select the required CI(s) at different levels of the view hierarchy. You can select multiple CIs by holding down the CTRL button and clicking the required CIs. To view and edit a CI's attributes, see "Displaying CI Attributes" on page 429.

➤ In the **Search** tab, search for CIs in existing views or in the entire CMDB. For details, see "Searching for a CI" on page 429.

➤ In the **Related CIs** tab, select a CI that is related to the selected CI in the **Topology View** tab. For details, see "Displaying Related CIs" on page 431.

Comparing CIs by View

Comparing CIs by view allows you to display the compared CIs and their child CIs that appear in the selected view.

To compare CIs by view:

1 In Topology View or View Explorer, select the two CIs you want to compare.

2 Right-click and select **Compare Compound** > **View** to display the selected CIs and their child CIs that appear in the view. The Compare compound CIs from view dialog box opens.
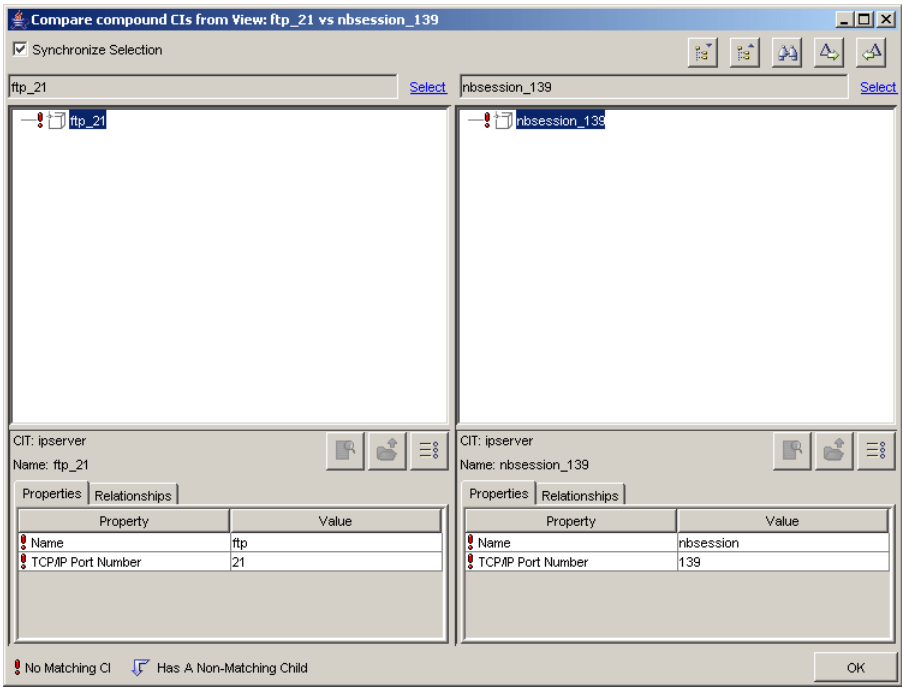
Understanding the Compare CIs Dialog Box

The Compare CIs dialog box is divided into two sections. Each side displays an hierarchical tree structure of the CIs you selected. When you select a CI on one side, the corresponding CI on the other side is automatically selected, provided the **Synchronize Selection** check box is selected. (For details, see "Enabling Synchronized Selection" on page 445.)

The following information is displayed under the tree when you select a CI:

CIT – The name of the CIT to which the CI belongs as defined in the CI Type Manager.

Name – The name of the CI as defined in the CI's label.

Understanding the Properties Tab

The **Properties** tab displays only the attributes of the CI whose values have been defined as **Comparable**. For more information about the **Comparable** option, see "Creating CITs" on page 384.

The **Properties** tab has the following two fields:

Field	Description
Property	Displays the attribute name for the selected CI which was marked as **Comparable** in the CI Type Manager.
Value	Displays the property value for the currently selected CI.

Expanding/Collapsing the Tree

This section describes how to expand and collapse the tree.

To expand the entire tree:

Select a CI and click the **Expand All** button.

To collapse the tree:

Select a CI and click the **Collapse All** button.

Finding Items Marked as Changed

This section describes how to find the items that have been marked as changed. For information on the change indicators, see "Understanding Change Indicators" on page 438.

To find the next occurrence of an item marked as changed:

Click the **Next Difference** button.

To find the previous occurrence of an item marked as changed:

Click the **Previous Difference** button.

Finding a CI by Its Label

You can search for a CI within the selected view whose label contains the searched text. You can specify a whole word or part of it.

To find the CI for which you are looking:

1 Click the **Find** button to open the Find CI by label dialog box.
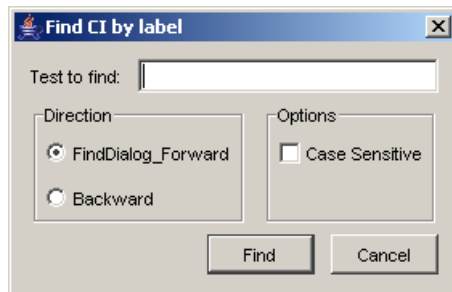
2 In the **Text to find** box, type the text you want to find.

3 In the **Direction** section, specify whether you want to do a forward or backward search.

4 In the **Options** section, select the **Case Sensitive** box if you want the matches to be case sensitive.

5 Click **Find**. If Mercury Application Mapping finds a match, it selects the CIs.

Understanding Change Indicators

This section describes the icons that indicate the kind of change that has occurred in the compared CIs. The indicator key appears at the bottom of the Compare compound CIs dialog box.

This section includes:

➤ "No Matching CI Icon" on page 439

➤ "Has a Non-Matching Child Icon" on page 439

No Matching CI Icon

The **No Matching CI** icon indicates that a CI on one side does not match any corresponding CI on the other side. The example below shows that the CI called **C** has the **No Matching CI** icon because of the change to the disk size (see the **Properties** tab underneath the tree structure on the right side).

Has a Non-Matching Child Icon

The **Has a Non-Matching Child** icon indicates that the CI has a match but one of its children does not have a match. The **disk** CI in the illustration above has a **Has a Non-Matching Child** icon due to a change to its child CI.

Viewing and Comparing Configuration Files

The Compare compound CIs dialog box enables you to view a CI's configuration file, and also compare the configuration files of two different CIs.

This section includes:

➤ "Viewing a CI's Configuration File" on page 440

➤ "Showing the Difference Between Two Configuration Files" on page 441

Viewing a CI's Configuration File

This section describes how to view a CI's configuration file.

To view a CI's configuration file:

 1 In the Compare compound CIs dialog box, select a CI of the CIT **configfile** from the tree.
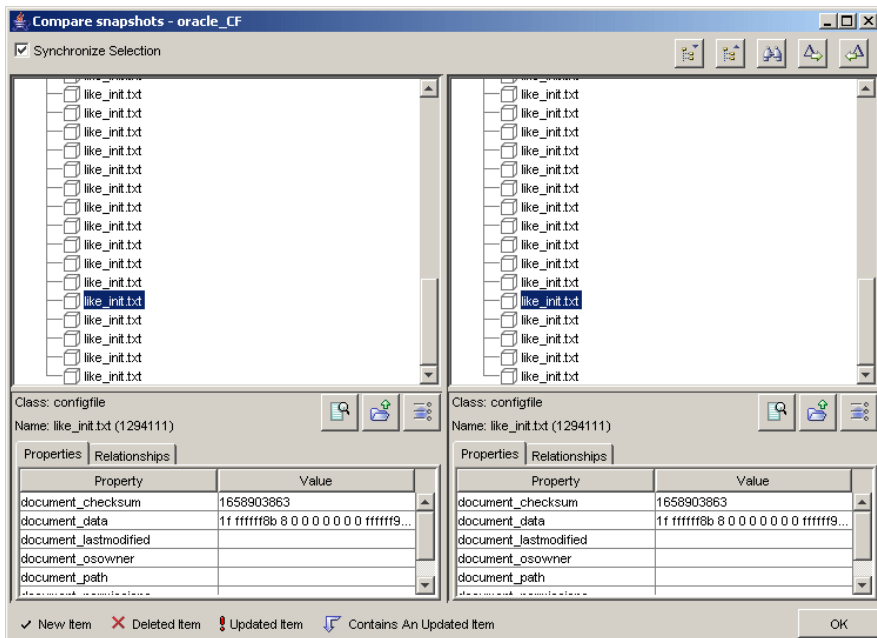
[icon] **2** Click the **Open** button to open the CI's configuration file.

Note: This button is only enabled if you have selected a CI of the CIT **configfile** from the tree.

Showing the Difference Between Two Configuration Files

This section describes how to show the difference between two configuration files.

To show the difference between the two configuration files:

1 In the Compare compound CIs dialog box, select a CI of the CIT **configfile** from the tree on each side.

[icon] **2** Click the **Show Difference** button.

Note: This button is only enabled if you have selected a CI of the CIT **configfile** from the tree.

The differences in the configuration files are marked in blue.

Displaying a CI's Attributes

You can display the selected CI's attributes.

To display a CI's attributes:

1 In the Compare compound CIs dialog box, select a CI from one of the trees.

2 Select the **Show CI Attributes** button to open the CI Attributes dialog box.
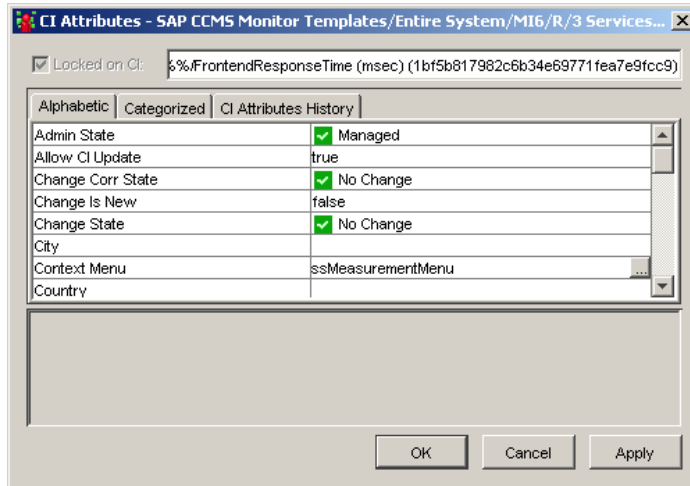
The CI Attributes dialog box enables you to view the attributes defined for the selected CI. The name of the selected CI appears at the top.

The CI Attributes dialog box has the following tabs:

➤ **Alphabetic** – Displays the CI's attributes in alphabetical order.

➤ **Categorized** – Displays the CI's attributes according to the CIT to which it belongs. Click the **Expand** button to display the CI attributes for that CIT.

➤ **CI Attributes History** – Displays a list of attributes in which changes have occurred.

Note: The CI Attributes History tab only displays the attributes whose values have been defined as **Change Monitored** in the Add Attributes dialog box in the CI Type Manager. For details, see "Creating CIT Methods" on page 391.

The CI Attributes History tab has the following fields:

Field	Description
Attribute	The name of the attribute.
Change Date	The date when the change occurred.
Changer	Indicates the cause of the change. For example, the field can contain the name of a user or a discovery pattern.
New Value	The new value of the attribute.

You can reduce the number of attributes that appear in the CIs Attributes History tab by selecting an attribute from the **Filter by** list.

Note: To view the updated information in the CI Attributes History tab, close the dialog box and then reopen it.

Enabling Synchronized Selection

Select the **Synchronize Selection** check box at the top-left corner of the Compare compound CIs dialog box to enable you to select a CI on one side and have its corresponding CI automatically selected on the other side.

The default for the **Synchronize Selection** check box is selected. Clearing this check box allows you to compare two different CIs by selecting one on each of the two trees. This allows you to compare two different CIs and see their different properties.

The nonmatching properties is marked by a **No Matching CI** icon in the **Properties** tab (see "No Matching CI Icon" on page 439).

26

Configuration File Support

This chapter describes how to view the content of configuration files and compare the differences between them. Configuration files use the **.xml**, **.txt** or **.ini** extensions.
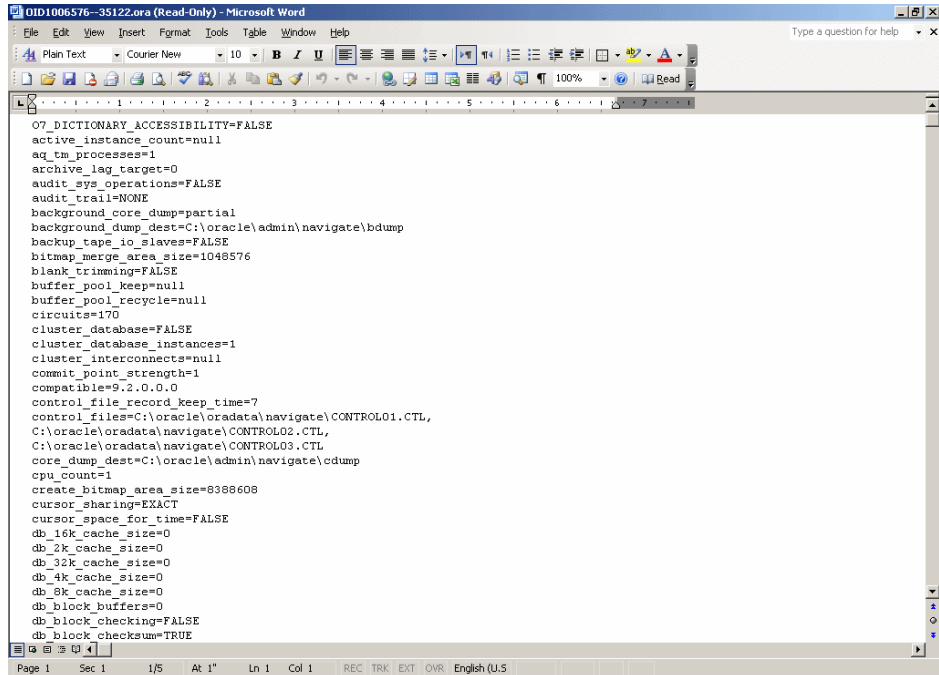
This chapter describes:	On page:
Viewing the Content of a Configuration File	448
Showing the Difference between Two Configuration Files	449
Showing the Differences for the Same Configuration File Over Time	450

Viewing the Content of a Configuration File

This section describes how to open a configuration file.

To open a configuration file:

1 Select a CI of the CIT **configfile**.

2 Right-click and select **Open** to view the content.

Showing the Difference between Two Configuration Files

This section describes how to compare the content of two configuration files.

To show the difference between the content of two configuration files:

1 Select two configuration files whose content you want to compare.

2 Right-click and select **Show Differences** to display the differences between the two configuration files.
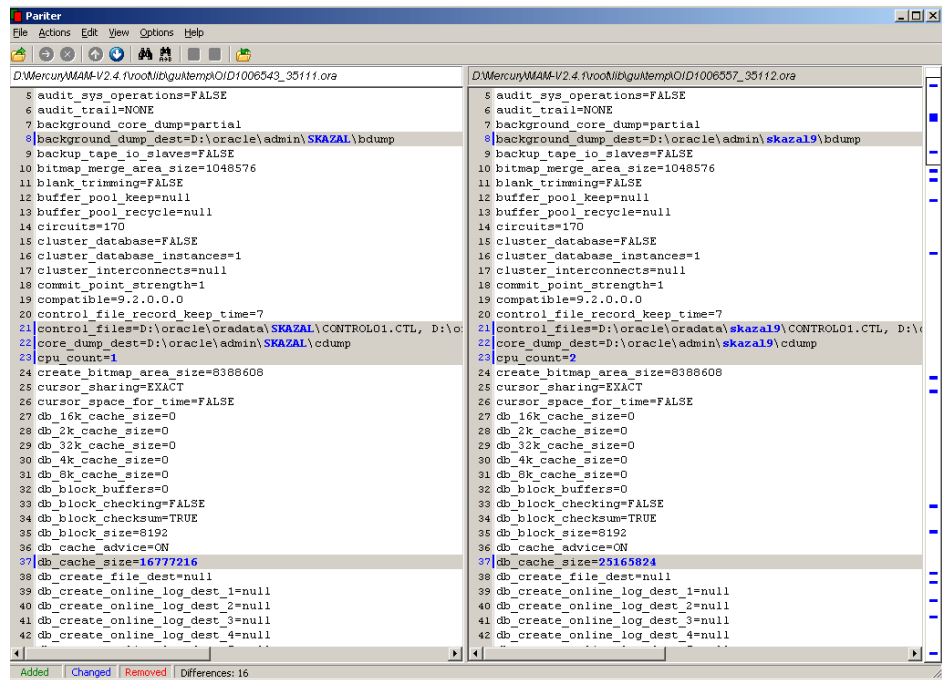
Showing the Differences for the Same Configuration File Over Time

This section describes how to display the differences for the content of a configuration file between two given dates.

1 Select the configuration file whose content you want to compare over time.

2 Right-click and select **Show History** to open the Date dialog box.
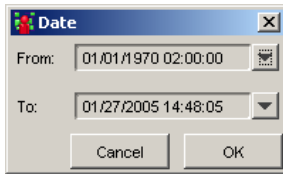
3 Select the dates during which you want to want to display changes that occurred in the file. To do this, do the following:

➤ Click the buttons to the right of the **From** and **To** fields to define the beginning and end of that time period respectively.
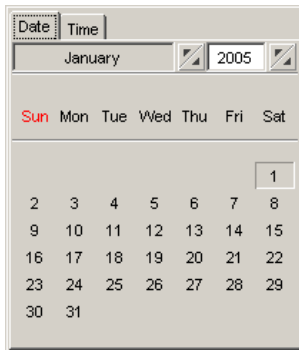
➤ Click the **Date** tab to select the month, date and year.

➤ Click the **Time** tab to select the required time.

➤ Click the button at the end of the **From** and **To** boxes to close the calendar.

➤ Click **OK** to open the Show History dialog box.
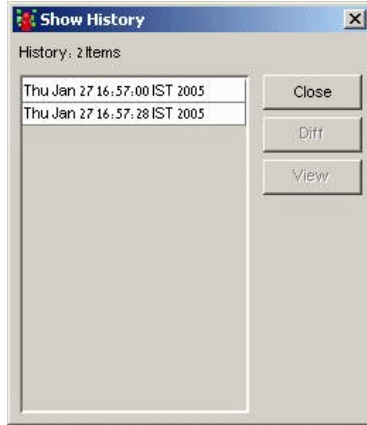
This dialog box lists all the times at which changes occurred in the configuration file.

4 From the **History** list, select the two items you want to compare.

5 Click **Difference** to display what the configuration file looked like at each time.

Note: The differences in the configuration file are marked in blue.

27

Impact Analysis

This chapter explains how Mercury Application Mapping enables you to assess the impact of infrastructure events.

About Evaluating the Impact of Infrastructure Events

Mercury Application Mapping enables you to simulate how changes you want to make will impact your system. Using the Ticket Manager, you create a Ticket, which includes the changes you want to make in the system. For example, a new software installation or upgrade. You then add the required CI(s) to the Ticket. The CIs you add to the Ticket represent the changes you want to make to the system. Mercury Application Mapping evaluates the impact the CIs associated with the Ticket will have on the system when these CIs are in downtime. You then generate a report displaying the CIs in the system, which will be impacted by the changes specified in the Ticket.

Setting Ticket Manager Configurations

The Ticket Manager can appear in different modes, depending on how you configure the following definitions in the **appilogConfig.properties** file, located in: \<**Mercury Application Mapping root directory**>**MAMServer\root\lib\server**.

➤ **impact.analysis.allow.ci.edit** – Determines whether or not you can:

 ➤ Modify the list of assets in a ticket.

 ➤ Add or remove CIs and save those changes.

 False = does not enable you to create new Tickets. False is the default.

 True = enables you to create new Tickets.

➤ **impact.analysis.read.only** – Determines whether or not you can create and modify ticket properties.

 False = does not enable you to modify Ticket properties.

 True = enables you to modify Ticket properties. True is the default.

Evaluating the Impact of Infrastructure Events

This section describes how to evaluate the impact of infrastructure events.

To evaluate the impact of infrastructure events:

 1 Click **Configuration Management** > **Impact Analysis** to open the Ticket Manager screen.

The Ticket Manager screen is divided into the following panes:

Pane	Description
Ticket List	Contains the following columns: ➤ A list of the Tickets that were created ➤ The report that was generated for that Ticket
Ticket Details	The Ticket's details.

 2 In the Ticket List pane, click the **Add New Ticket** button to create a new Ticket. The **Create** button appears in the Ticket Details pane.

 3 In the Ticket Details pane, define the Ticket's details as required and click **Create**. The new Ticket appears in the Ticket List pane.

 4 To remove an existing ticket from the Ticket List pane, select the ticket you want to remove and click the **Delete Ticket** button.

 5 Click **Asset Selection** in the Ticket Details pane to open a window that allows you to select the CI(s) you want associated with the selected Ticket.

 6 In the **Tree** tab, select the check box to the left of the required CI(s) and use the right arrow to add CIs to the Ticket Details pane.

 7 To remove a CI from a Ticket, select the check box to the left of the required CI(s) in the Ticket Details pane and click the left arrow.

 8 To find a CI using the Search tab, do the following:

 ➤ Click the **Search** tab.

 ➤ Enter the text of the label of the CI for which you want to search.

 ➤ Click **Go**. The Search tab returns all CI occurrences whose labels contain the text entered or part of it.

 ➤ Select the check box to the left of the required CI(s) and use the right arrow to add CIs to the Ticket Details pane.

 ➤ To remove CIs from the Ticket Details pane, select the the check box to the left of the required CI(s) in the Ticket Details pane and click the left arrow.

 9 Click the **Show Impact** button to generate an Impact Analysis report. This report displays a list of CIs that are impacted in the system as a result of the simulated changes.

 The Impact Analysis report displays the results as follows:

 ➤ The CIs in the **Impacted Services** tab are grouped according to the hierarchy of the folders and views in the Explorer Tree pane.

 ➤ The CIs in the **Impacted Resources** tab are sorted by CIT.

Modifying a Ticket's Details

This section describes how to modify an existing Ticket.

To modify an existing Ticket:

1 In the Ticket List pane, select the Ticket you want to modify.

2 In the Ticket Details pane, make the required changes.

Click **Save**.

Part XI

Mercury Application Mapping WebView

28

Introducing Mercury Application Mapping WebView

This chapter introduces Mercury Application Mapping WebView, which is a Web-based, management console designed to enable business service management.

This chapter describes:	On page:
About Introducing Mercury Application Mapping WebView	459
Accessing Mercury Application Mapping WebView	460
Mercury Application Mapping WebView Main Window	462
Navigating Mercury Application Mapping WebView	469
Mercury Application Mapping WebView Buttons	469

About Introducing Mercury Application Mapping WebView

You use Mercury Application Mapping WebView to:

➤ manage business services – enables you to view predefined maps using a Web interface

➤ view event information – enables events that are triggered in Mercury Application Mapping to be seen in Mercury Application Mapping WebView

➤ perform Mercury Application Mapping actions – enables you to perform actions on CIs such as problem isolation and show impact

---

Note: Mercury Application Mapping WebView does not have any administrative functionalities, such as building TQLs or creating correlation rules. For details on administering Mercury Application Mapping, refer to the *Mercury Application Mapping Administration Guide*.

Accessing Mercury Application Mapping WebView

You access Mercury Application Mapping WebView using a supported Web browser, from any computer with a network connection (intranet or Internet) to the Mercury Application Mapping servers. The level of access granted depends on your permissions. For details on granting user permissions, see "User and Role Management" in *Mercury Application Mapping Administration Guide*.

For details on Web browser requirements, as well as minimum requirements to successfully view Mercury Application Mapping, see "Hardware and Software Requirements" in *Mercury Application Mapping Installation Guide*.

This section includes:

➤ "Logging In" on page 460

➤ "Logging Out" on page 461

Logging In

You log into Mercury Application Mapping WebView from the login page.

To access the login page and log in:

1 In the Web browser, enter the URL **http://<ipaddress:portnumber>/webapp/**, where **ipaddress** is the IP address of the Mercury Application Mapping Server and **portnumber** is the port number of the J2F server. The default value is **8080**.

2 Enter the login parameters (login name and password) of a user defined in the Mercury Application Mapping system, and click **Login**. After logging in, the user name appears at the top right, under the top menu bar.

Initial access can be gained using the default superuser login parameters: login name=**Administrator**, password=**Administrator**. It is recommended that the system administrator change this password immediately to prevent unauthorized entry. For details on changing the password, see "Changing a User's Password" in the *Mercury Application Mapping Administration Guide*.

For details on creating users in the Mercury Application Mapping system, see "Defining a User" in the *Mercury Application Mapping Administration Guide*.

Logging Out

When you have completed your session, it is recommended that you log out of the Web site to prevent unauthorized entry.

To log out:

Click **LOGOUT** at the top right corner of the page.

Mercury Application Mapping WebView Main Window

The Mercury Application Mapping WebView main window includes the following panes: .

Refresh button

Search pane   Automatic Refresh

Explorer Tree pane   Map Detail pane

This section includes:

➤ "Search Pane" on page 463

➤ "Explorer Tree Pane" on page 463

➤ "Map Detail Pane" on page 464

Search Pane

You can search for CIs and views that are displayed in the Explorer Tree using the Search pane (for details, see "Searching for a View or a CI" on page 473).

Explorer Tree Pane

The Explorer Tree pane displays a hierarchical tree structure of the CIs and the folders in which they are located.

The tree includes folder, view, and CI icons and displays four levels:

➤ **Root** – contains all folders and views in the tree hierarchy. This is the highest level.

➤ **Folder** – contains a specific view or views.

➤ **View** – displays the view itself.

➤ **CI** – displays the CI itself.

Select a level in the Explorer Tree pane to display all CIs related to the selected layer in the Map Detail pane. For example, if you select a folder in the Explorer Tree pane, the Map Detail pane displays the views that the selected folder contains. If you select a view in the Explorer Tree pane, the Map Detail pane displays its CIs.

Note: The Explorer Tree is automatically updated to reflect any change that occurs in the View Explorer in Mercury Application Mapping.

Map Detail Pane

The Map Detail pane displays the results of the view that is currently selected in the Explorer Tree pane. For details on the Explorer Tree pane, see above.

The Map Detail pane includes the following tabs:

➤ Layer tab (for details, see below)

➤ Properties tab (for details, see "Properties Tab" on page 466)

➤ Events tab (for details, see "Events Tab" on page 467)

➤ Map tab (for details, see "Map Tab" on page 468)

Layer Tab

The layer tab displays the topology and status data for the currently selected CI.
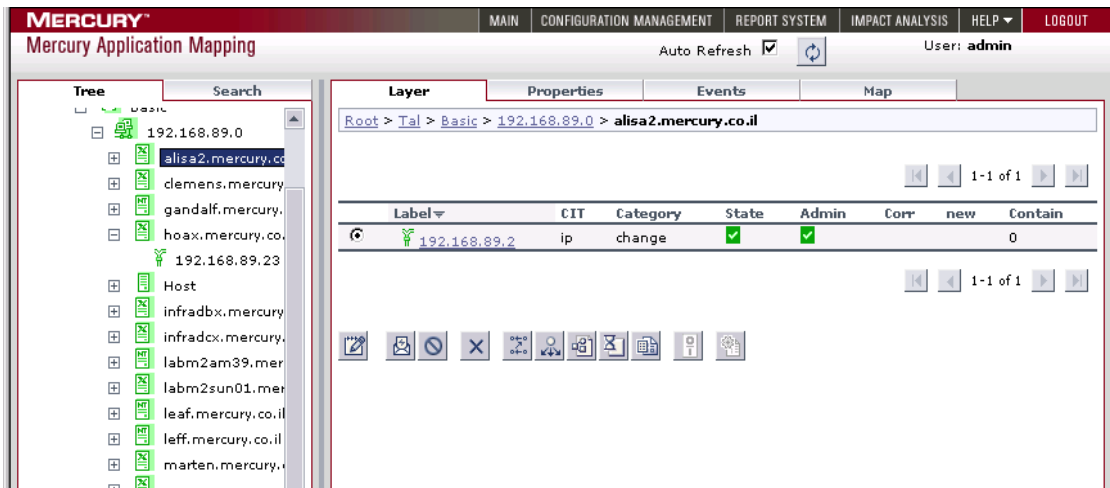
The CIs represented in the topology map include topology data as well as basic attributes. This topology data is based on, and produced according to, the way these CIs are located and defined in Mercury Application Mapping views.

The **Layer** tab displays the following details for the selected CI:

➤ **Label** – The label of the selected CI.

➤ **CIT** – The CIT of the selected CI as defined in the CIT Model.

➤ **Category** – The category that is defined for the selected CI in the Category Manager dialog box.

Note: If more than one category is defined for the CI, the category that has the highest priority as defined in the **Category Priority** list is displayed.

➤ **State** – The operational state of the CI, as adjusted for the weight accorded to its significance in the system. The significance weight is set in the Status Factor tab. For details, see "Specifying the CI Status Factor" in *View Manager Administration*.

➤ **Admin** – The administrative state of the selected CI. Any color other than green indicates that the event information coming from the selected CI is not reliable because the CI is disabled, undergoing testing, or is otherwise unmanaged. The Admin State of a CI does not propagate upwards to CIs at higher levels.

➤ **Corr** – A state that is defined in a correlation rule, and assigned to an affected CI once its root cause CI fulfills the correlation condition and a correlation event occurs. For details on correlation rules, see Chapter 15, "Defining Correlation Rules."

➤ **New** – An exclamation mark in this field indicates that the CI's active event has not been acknowledged.

➤ **Contain** – The number of children that exist at all map levels below the selected level.

Note: You can sort the contents in each column of the Layer tab by clicking the sort button in the column header. To change the sort order, click the sort button in the column header again.

Properties Tab

The **Properties** tab displays the CI details according to the format defined in its CIT report. It is enabled at the CI level.

The **Properties** tab displays the details of the CIT report definition of the selected CI. The details that appear vary, as they are specific to the CIT for which you defined the report.

Events Tab

You use the **Events** tab to view event information. The content of the Events tab depends on which button is clicked (**Compound Events**, **Event History**, **Toggle QuickFilter** is **On/Off**, or **Acknowledge All**). For details, see "Mercury Application Mapping WebView Buttons" on page 469.

The Events tab contains the following fields:

➤ **Severity** – The severity of the event (Normal, Warning, Minor, Major, Critical). Each severity level is displayed in a different color.

➤ **Category** – The category for which the event is sent.

➤ **Class** – The CIT name of the event's CI.

➤ **Label** – The inclusive label of the CI related to the event. Includes other CIs that identify the selected CI. For example, a disk's label includes the disk name and its host name.

➤ **Msg** – A textual description of the event.

➤ **Ack** – Indicates whether the event has been acknowledged or not.

➤ **Discovery Time** – The date and time at which the event is triggered.

➤ **Corr** – A state that is defined in a correlation rule, and assigned to an affected CI once its root cause CI fulfills the correlation condition and a correlation event occurs. For details, see Chapter 9, "Using the Topology Map."

➤ **Origin** – Indicates the origin of the event, for example, a correlation rule, a user, or a discovery pattern.

Note: You can sort the contents in each column of the Events tab by clicking the sort button in the column header. To change the sort order, click the sort button in the column header again.

Map Tab

You use the Map tab to view multi-level maps based on a selected view. You can navigate between different map levels (from the top level down to the CI level) to focus on specific areas of a map. For a description of the icons in the Map tab, see "Toolbar" on page 34.
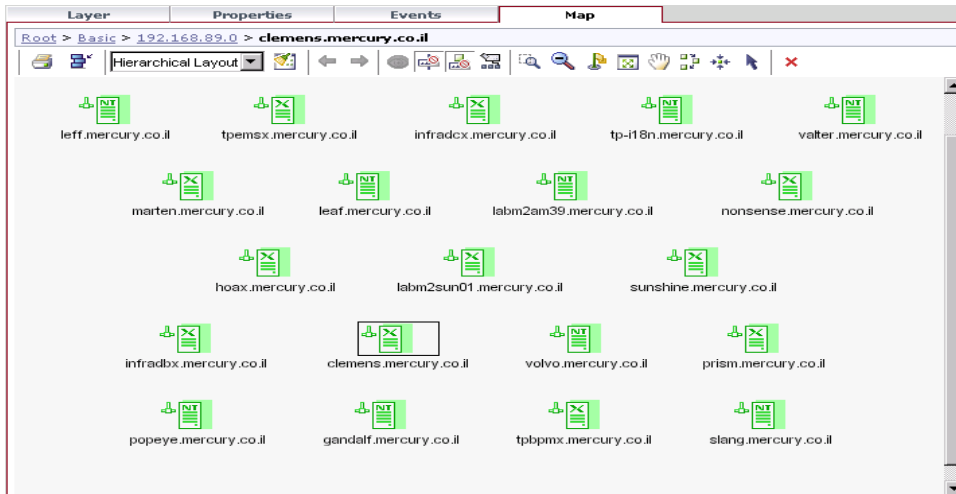
Navigating Mercury Application Mapping WebView

Mercury Application Mapping WebView runs in a Web browser. You move around Mercury Application Mapping WebView using the following navigation functions:

➤ **Menus** – enable navigation to various contexts within Mercury Application Mapping WebView. The menus accessed from the top menu bar provide navigation to the views, reports, and help resources. The top menu bar also includes a link to the Logout button.

➤ **Tabs** – enable navigation to various contexts within a particular function, for example to different information about a CI.

➤ **Breadcrumbs** – You can return to previous pages within a multi-level context by clicking the appropriate page level.

➤ **Drill down** – CIs that include child CIs are displayed as relationships. You drill down through the CIs by clicking the relationship.

Mercury Application Mapping WebView Buttons

Mercury Application Mapping WebView includes the following buttons:

Button	Description	Displayed in the Following Tab(s)
	Compound Events – Toggles between showing or hiding the active events of the selected CI and of its children. For details, see "Viewing Event History" on page 480.	➤ Event
	Event History – Toggles between showing or hiding the log events of the currently selected CI and its children. The **Compound Events** must be toggled **on** to show log events.	➤ Event

Button	Description	Displayed in the Following Tab(s)
	Toggle QuickFilter is On – Toggles between showing or hiding the events that have been filtered in the Quick Filter dialog box. For details, see "Filtering the Events You Want to Display" on page 481. **Note:** This button only toggles if an event filter is defined for the CI.	➤ Event
	Toggle QuickFilter is Off – Includes the events that are filtered in the Quick Filter dialog box.	➤ Event
	Next Page – Go to next page.	➤ Layer ➤ Event
	Last Page – Go to last page.	➤ Layer ➤ Event
	Previous Page – Go to previous page.	➤ Layer ➤ Event
	First Page – Go to first page.	➤ Layer ➤ Event
	Refresh – Manually updates the information with the latest topology and events data. Use this button if you do not want to wait for the periodical refresh.	➤ All tabs
Auto Refresh ☑	**Auto Refresh** – Automatically updates the information with the latest topology and events data. When the button is checked, a refresh is periodically done. When the button is not checked, a refresh is not periodically done.	➤ All tabs
	Edit – for details, see "Editing a Note and Attaching It to a CI" on page 476	➤ Layer
	Send Event – for details, see "Sending an Event" on page 252	➤ Layer

Button	Description	Displayed in the Following Tab(s)
	Clear Events – for details, see "Clearing All Events" on page 252	➤ Layer
	Delete Object – for details, see "Deleting a CI" on page 476	➤ Layer
	Problem Isolation – for details, see "Isolating a Problem for a CI" on page 477	➤ Layer
	Show Impact – for details, see "Showing the Correlation Impact of a CI" on page 476	➤ Layer
	Show Related CIs Report – For details, see .	➤ Layer
	Acknowledge All – For details, see "Acknowledging All Events Related to a Selected CI" on page 481.	➤ Event
	Object History – Lists the history of changes to CI attributes that have been saved in the history database. For details, see "Viewing and Editing CI Attributes" on page 201.	➤ Layer
	Change Report – Lists the history of changes to a specific CI and its subgroup in the database. For details, see "Generating a Change Report" on page 268.	➤ Layer
	Activate/Deactivate View – Toggles between activating and deactivating the view state. For details, see "Deactivating and Activating Views" on page 196.	➤ Layer
	View Object Documents – Opens a file that has any type of data (xml, txt, doc, and so forth). The data attribute is **byte[]** and its qualifier is **viewable**.	➤ Layer

29

Using Mercury Application Mapping WebView

This chapter explains how to use Mercury Application Mapping WebView.

For a description of the Mercury Application Mapping WebView window, see "Mercury Application Mapping WebView Main Window" on page 462.

Searching for a View or a CI

The Search pane enables you to search for views or CIs that are displayed in the Explorer Tree.

To search for a view or a CI:

1 In the **Search** box, enter the name of the view or CI or part of it.

2 Click **Go**. The search results appear in the Search pane displaying all the instances of the text that appear in Mercury Application Mapping.

Note: Mercury Application Mapping WebView does not consider case when differentiating between strings. That is, Server is considered the same as server.

Displaying All CIs Related to a Selected Layer

Select a level in the Explorer Tree pane to display all CIs related to the selected layer in the Map Detail pane. For example, if you select a folder in the Explorer Tree pane, the Map Detail pane displays the views that the selected folder contains.

If you select a view in the Explorer Tree pane, the Map Detail pane displays its CIs. You can see the view's nodes and CIs only in the Map Detail pane.

Note: Every time you navigate to a different level, the Orientation Bar displays your current location in the map hierarchy (for details, see "Navigating Mercury Application Mapping WebView" on page 469).

The Explorer Tree is automatically updated to reflect any change that occurs in the View Explorer in Mercury Application Mapping.

Viewing Details of a CI in the Layer or Events Tab

You can view a selected CI's properties and you can perform actions, such as showing the root-cause, attaching a note, and so forth.

To view details of a CI in the Layer or Events tab:

Select the CI in the Explorer Tree pane. The Map Detail pane displays the CI properties. Buttons are located at the bottom of the Map Detail pane. For details about button functions, see "Mercury Application Mapping WebView Buttons" on page 469.

You can perform the following actions on the CI by clicking the desired button:

➤ **Edit Note** (for details, see "Editing a Note and Attaching It to a CI" on page 476)

➤ **Send Event** (for details, see "Sending an Event" on page 477)

➤ **Clear Events** (for details, see "Clearing Events from a CI" on page 478)

➤ **Delete Object** (for details, see "Deleting a CI" on page 476)

➤ **Problem Isolation** (for details, see "Isolating a Problem for a CI" on page 477)

➤ **Show Impact** (for details, see "Showing the Correlation Impact of a CI" on page 476)

➤ **Show Related CIs Report** (for details, see "Isolating a Problem for a CI" on page 477)

➤ **Object History** (for details, see "Viewing and Editing CI Attributes" on page 201)

➤ **Change Report** (for details, see "Generating a Change Report" on page 268)

➤ **Activate / Deactivate View** (for details, see "Deactivating and Activating Views" on page 196)

➤ **View object documents**

Editing a Note and Attaching It to a CI

You can edit a note and attach it to a node. You can use these notes as reminders to yourself or to provide feedback for others.

To edit a note:

1 In the Map Detail pane, click the **Layer** tab.

2 In the **Layer** tab, select the CI to which you want to add a note.

3 Click **Edit Note** to open the Edit Note dialog box.

4 Type the required text in the Note box.

5 Click **Submit** to save the note.

Deleting a CI

You can delete a CI.

To delete a CI:

1 In the Map Detail pane, click the **Layer** tab.

2 In the **Layer** tab, select the CI to delete.

3 Click **Delete Object**. A confirmation message is displayed.

4 Click **Submit** to delete the CI or **Cancel** to return to the CI without deleting it.

Showing the Correlation Impact of a CI

You can select a CI which is defined by a correlation rule as a root-cause CI, and display all CIs that are affected by it in a separate window. For additional information, see Chapter 9, "Using the Topology Map."

To show a correlation impact on a CI:

1 In the Map Detail pane, click the **Layer** tab.

2 In the **Layer** tab, select a root-cause CI.

3 Click **Show Impact**. A Show Proximate Impact window opens.

Isolating a Problem for a CI

You can select a CI that is affected by a correlation rule and display the root-cause CIs that affect it.

To isolate a problem for a CI:

1 In the Map Detail pane, click the **Layer** tab.

2 Select a CI (click next to it in the table row) that is affected by a correlation rule.

3 Click **Problem Isolation.** A window opens displaying the CI affected by the correlation rule and all root-cause CIs that affect it.

Sending an Event

You can define the category that will be affected by the event to be sent.

To send an event:

1 In the Map Detail pane, click the **Layer** tab.

2 In the **Layer** tab, select the CI to which you want to send an event.

3 Click **Send Event** to open the Send Event dialog box.

From the category list, select the category that is affected by the event. The categories **operation, change**, and **test** are system defaults. Other categories is the list are dynamically defined in the Category Manager.

4 Select the required severity.

5 Click **Submit**. The button defined for the severity you selected is displayed.

Clearing Events from a CI

You can delete a CI's events, either those that belong to a certain category or to all categories.

To clear events from a CI:

1 In the **Layer** tab, select the CI from which you want to clear events.

2 Click **Clear Events** to open the Clear Events dialog box.

3 Select a category in the **Category** list to delete events belonging to a specific category. Note that the list contains categories that are either the system's default categories or are defined in the Enumeration Manager.

4 Select **All Categories** to delete events belonging to all categories

---

Note: When clearing all events for a specific category, the CI may still be affected by other category events. Only if you choose to delete the events for all category types, does the CI's state return to normal.

---

5 Click **Submit** to save the settings.

Opening a Map of the Currently Selected View

You use the Map tab to view multi-level maps based on a selected view. You can navigate between different map levels (from the top level down to the CI level) to focus on specific areas of a map.

To open a map of the currently selected view:

1 In the Explorer Tree pane, select the required view.

2 In the **Map** tab, the map of the currently selected view is displayed.
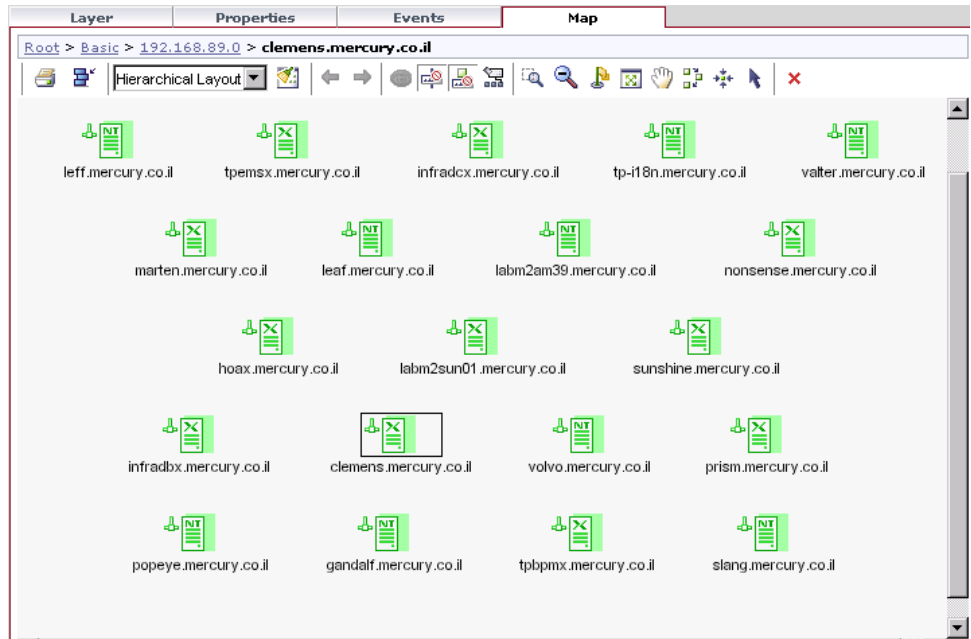
3 Right-click a CI to open a menu that enables you to view more information. For details on the menu options, see "CI Shortcut Menu" on page 170. For details about the buttons that appear in the Map tab, see "Toolbar" on page 34.

Evaluating the Impact of Infrastructure Events

Mercury Application Mapping WebView enables you to simulate how changes you want to make will impact your system. Using the Ticket Manager, you create a Ticket, which includes the changes you want to make in the system. Click **Show Impact** at the top to open the Ticket Manager. For more details, see "Impact Analysis" on page 453.

Viewing Event Information

You can view information about events as follows:

➤ Display the event history of a selected CI (for details, see "Viewing Event History" on page 480).

➤ Acknowledge all events relating to a selected CI. (for details, see "Acknowledging All Events Related to a Selected CI" on page 481).

➤ Display only specific events in the Events tab for a specific CI (for details, see "Filtering the Events You Want to Display" on page 481).

Viewing Event History

You can view the event history of the selected CI. Mercury Application Mapping WebView lists the raw and active events that occur in relation to the selected CI and its children.

In the Events tab, you use the compound and history buttons to toggle between four different states:

State	When This Button...	Is Pressed	Is Not Pressed	This is Displayed
State 1	Compound		√	Active events of the currently selected CI.
	History		√	
State 2	Compound	√		Active events of the currently selected CI and all its children.
	History		√	

State	When This Button...	Is Pressed	Is Not Pressed	This is Displayed
State 3	Compound		√	Raw events of the currently selected CI.
	History	√		
State 4	Compound	√		Raw events of the currently selected CI and all its children.
	History	√		

Acknowledging All Events Related to a Selected CI

You can acknowledge all events relating to a selected CI.

To acknowledge all events related to a selected CI:

1 In the Map Detail pane, click the **Events** tab.

2 In the **Events** tab, select a CI to which events have been sent.

3 Click **Acknowledge All**. A mark appears in the **Ack** column.

Note: Acknowledge All is only for active events.

Filtering the Events You Want to Display

You use the QuickFilter dialog box to include only events you want to display in the Events tab for the selected CI.

To filter the events you want to display:

1 In the Explorer Tree pane, select a CI to which an event has been sent.

2 In the Map Detail pane, click the **Events** tab.

3 Click **QuickFilter Settings** to open the QuickFilter dialog box.

4 To display events that were sent during a specified period of time:

➤ Click the date to the right of the **From** and **To** boxes to define the beginning and end of the time period.

➤ In the **Time** box, enter the required time.

➤ Click a date on the calendar. The times and dates you select are displayed in the **From** and **To** boxes.

➤ Click **OK**.

5 Select the required severity levels for the different categories.

6 From the **isAck** list, select an option:

True – display CIs whose events have been acknowledged

False – display CIs whose events have not been acknowledged

None – display CIs whose events are either in an acknowledged or unacknowledged state

7 From the **isCorr** list, select an option:

True – display CIs whose events are the result of a correlation rule

False – display CIs whose events are not the result of a correlation rule

None – display CIs whose events are either the result of a correlation rule or not the result of a correlation rule

---

Note: Use any necessary combination of parameters you need to correctly define the filter conditions.

---

8 From the **Sort By** list, select a category.

9 Click **OK** to save the settings you have defined. You are returned to the Events tab where only events that match the filter conditions are displayed.

10 Toggle the **Toggle QuickFilter** button **on** to display the filtered events.

30

Generating Reports in Mercury Application Mapping WebView

This chapter describes how to generate reports in Mercury Application Mapping WebView.

Understanding Reports in Mercury Application Mapping WebView

You can generate reports that are defined in Mercury Application Mapping and view them in Mercury Application Mapping WebView.

➤ **Asset Report** (for details, see "Generating Asset Reports" on page 485)

➤ **Host Dependency Report** (for details, see "Generating Host Dependency Reports" on page 487)

➤ **Gold Master Report** (for details, see "Generating Gold Master Reports" on page 488)

➤ **Compare Compound** (for details, see "Comparing Compound CIs" on page 491)

➤ **System Reports** (for details, see "Generating System Reports" on page 492)

➤ **Compare Snapshots** (for details, see "Comparing Shapshot Reports" on page 494)

Generating Asset Reports

You can generate a report that lists all the CIs in a selected view and their attribute values. For information on Asset Reports, see "Generating an Asset Report" on page 258.

To generate an asset report:

1 Click the **Configuration Management** tab at the top.

2 Click the **Asset Report** tab.

Note: The report contains only the attributes whose **Asset Data** check box was selected in the **Attributes** tab in the CI Type Manager. For more details, see "Creating CITs" on page 384.

> ➤ Select the **View only** check box to view only the attribute values of the CIs in the selected view.

> ➤ Clear the **View only** check box to view the attribute values of the CIs in the entire CMDB. The default is not selected.

3 In the Explorer Tree pane at the left, select the view for which you want to create an Asset report.

4 If you want to view a report for another view, select the required view and click the **Generate** button.

5 To print the report, do the following:

> ➤ In the top-right corner of the tab, click the **Printer Friendly** buttonto display the report in a separate window.

> ➤ Click the **Printer Friendly** button to print the report.

6 To view the report in Excel format, click the **Export to Excel** button in the top-right corner of the tab.

Generating Host Dependency Reports

You can generate a report displaying the dependencies between hosts. For details, see "Generating a Dependency Report" on page 260.

To generate a Host Dependency report:

1 Click the **Configuration Management** tab at the top.

2 Click the **Host Dependency Report** tab.

3 In the left pane, select the view for which you want to create a Host Dependency report.

➤ Select the **View only** check box to view only the dependencies between hosts in the selected view.

➤ Clear the **View only** check box to view the dependencies between hosts in the entire CMDB. The default is not selected.

4 If you want to view a report for another view, select the required view and click the **Generate** button.

 5 To print the report, do the following:

   ➤ In the top-right corner of the **Host Dependency** tab, click the **Printer Friendly** button to display the report in a separate window.

   ➤ Click the **Printer Friendly** button to print the report.

 6 To view the report in Excel format, click the **Export to Excel** button in the top-right corner of the tab.

Generating Gold Master Reports

Mercury Application Mapping WebView version enables you to generate a report that compares the configuration of a Gold Master CI to other CIs of the same CIT. You first perform the Gold Master comparison in the Mercury Application Mapping main window and then generate the report in Mercury Application Mapping WebView.

To generate a Gold Master report in Mercury Application Mapping WebView:

 1 Compare the configuration of CIs to a Gold Master template in the Mercury Application Mapping main window as described in "Comparing the Configuration of CIs to a Gold Master Template" on page 424.

 2 In Mercury Application Mapping WebView, click **Configuration Management** at the top.

3 Click the **Goldmaster Report** tab.

The following table describes the panes contained in the Gold Master Report tab:

Pane	Description
Left pane	Displays a list of all the Gold Master reports created in Mercury Application Mapping.
Goldmaster Report Details	Has the following fields: ➤ **Name** – the report name ➤ **Description** – the report description
Goldmaster Report Results	Displays the results of reports that are scheduled to run on a periodic basis in the Scheduled Actions Manager using the **Generate and Save a System Report** action. For details on using the Scheduled Actions Manager, see the *Mercury Application Mapping Administration Guide*. The Report Results pane contains the following fields: ➤ **Name** – the name of the report ➤ **Description** – a description of the report ➤ **Date** – the creation date of the scheduled report ➤ **View** – the reports that were generated at the time intervals specified in the scheduled report definition Note: To view the results of a report that was generated at specified time intervals, see "Viewing a Gold Master Report at Specific Intervals" on page 491.

 4 In the left pane, select the required Gold Master report.

 5 Click the **Generate Goldmaster Report** button to create the report. The Gold Master Report is displayed in a separate window. For a description of the fields a Gold Master Report has, see "Comparing the Configuration of CIs to a Gold Master Template" on page 424.

 6 To print the report, click the **Printer Friendly** button in the top-left corner of the report.

7 To view the report in Excel format, click the **Export to Excel** button in the top-left corner of the tab.

Viewing a Gold Master Report at Specific Intervals

You can view the results of a report that was generated at the time intervals specified in the scheduled report definition.

Note: If the report is not defined to be generated periodically, the Report Results pane is empty.

To view a Gold Master created at a specific time interval:

In the Report Results pane, click the **View Report** button in the **View** column of the report you want to generate. The Gold Master report is displayed.

Comparing Compound CIs

Mercury Application Mapping WebView version enables you to compare two CIs that are linked by a **container** relationship and see the hierarchical differences between them.

To compare compound CIs:

1 Click the **Configuration Management** tab at the top.

2 Click the **Compare Compound** tab. For details on how to compare compound CIs, see "Comparing Compound CIs" on page 432.

Generating System Reports

System reports include information that is defined in Mercury Application Mapping. For details, see Chapter 20, "Creating System Report Templates".

 1 Click the **Report System** tab at the top.
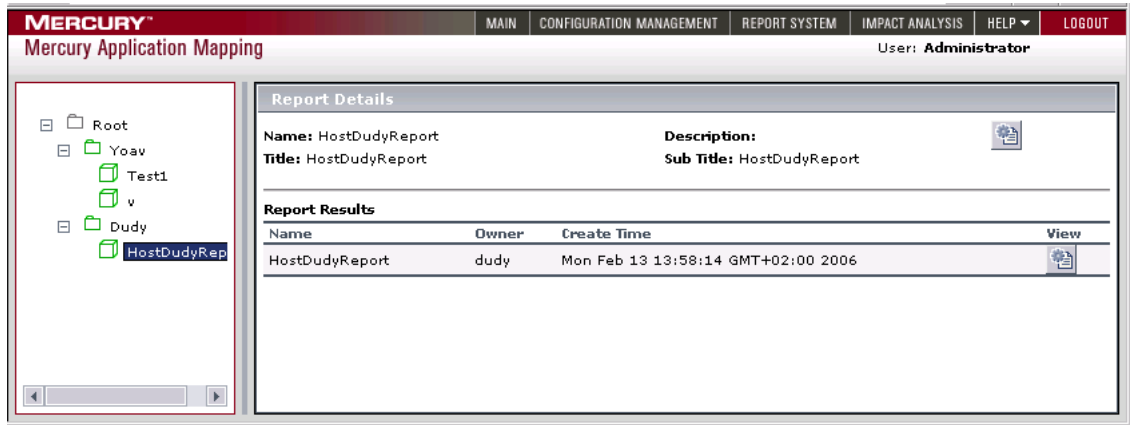
The following table describes the panes contained in the Report System tab:

Pane	Description
Left pane	Displays a hierarchical tree structure of the reports you defined in Mercury Application Mapping.
Report Details	Has the following fields: ➤ **Name** – The report name. ➤ **Description** – The report description. ➤ **Title** – The title of the report that appears at the head of the report. ➤ **Sub Title** – The subtitle of the report that appears underneath the title at the head of the report.
Report Results	Displays the results of reports that are scheduled to run on a periodic basis in the Scheduled Actions Manager using the **Generate and Save a System Report** action. For details on using the Scheduled Actions Manager, see the *Mercury Application Mapping Administration Guide*. The Report Results pane contains the following fields: ➤ **Name** – the name of the report ➤ **Owner**– the name of the user that ran the report ➤ **Create Time**– the time the report is created ➤ **View** – the reports that were generated at the time intervals specified in the scheduled report definition **Note:** To view the results of a report that was generated at specified time intervals, see "Generating a System Report at Specific Intervals" on page 494.

2 In the pane to the left, select the required report.

3 Click the **Generate Report** button in the Report Details pane to create the report. The report appears in a new window.

4 To print the report, click the **Printer Friendly** button in the top-left corner of the report.

5 To view the report in Excel format, click the **Export to Excel** button in the top-left corner of the tab.

Generating a System Report at Specific Intervals

You can also view the results of a report that was generated at the time intervals specified in the scheduled report definition.

Note: If the report is not defined to be generated periodically, the Report Results pane is empty.

To generate a system report created at a specific time interval:

In the Report Results pane, click the **View Report** button in the **View** column of the report you want to generate. The system report is displayed.

Comparing Shapshot Reports

Mercury Application Mapping allows you to compare two snapshots of a specific view taken at different times. This feature enables you to visualize the differences between the views by comparing their statuses at the times the snapshots were taken. For details on how to compare snapshots, see "Comparing Snapshots Taken at Different Times" on page 408.

To compare different snapshots of the same view:

1 Click the **Configuration Management** tab at the top.

2 Click the **Compare Snapshots** tab.
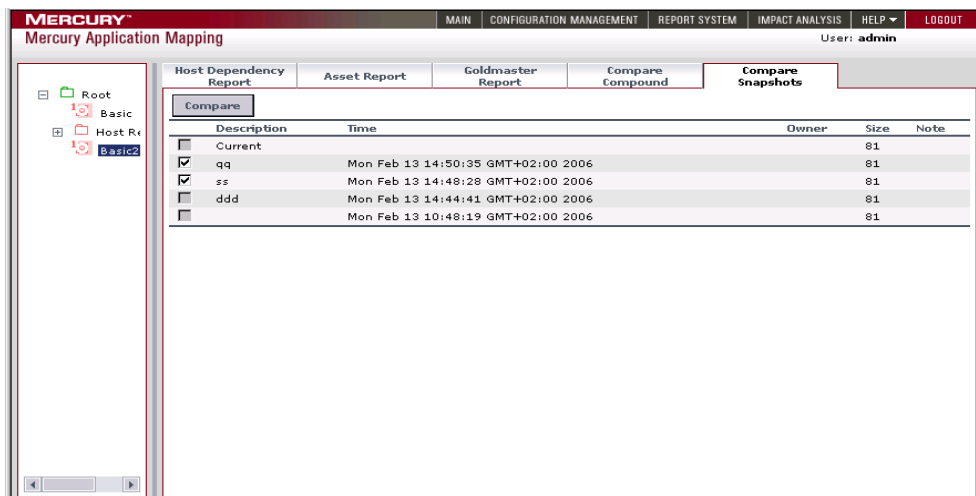
3 Select two snapshots for comparison and click **Compare.** A window opens with a side-by-side comparison of the selected snapshots.
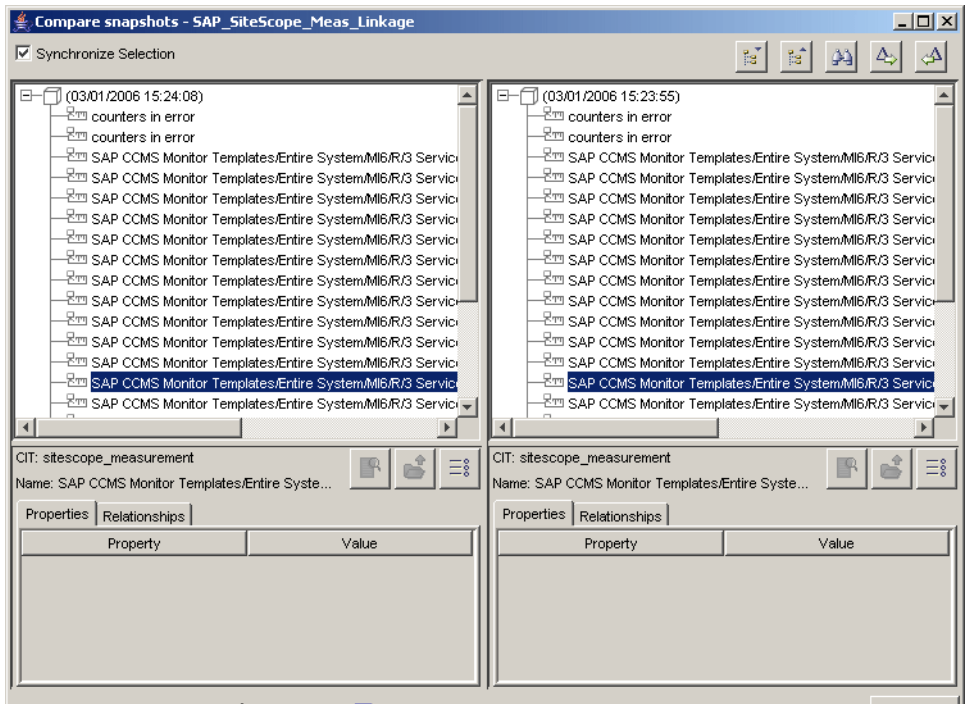
For details in using the Compare snapshots dialog box, see "Understanding the Compare Snapshots Dialog Box" on page 410.

Part XII

Appendixes

A

CI Type Descriptions

This section provides additional information about CI Types, permitted relationships between CIs, and key attributes of CITs. .

This chapter describes:	On page:
Diagram of CI Type Host and Permitted Relationships	499
Diagram of CITs Whose Parent is System	501
List of Parent, Description, and Key Attributes of all CI Types	501
Permitted Relationships between CI Types	532

Diagram of CI Type Host and Permitted Relationships

A CI can represent hardware, software, services, business processes, and so forth. CIs with similar properties are grouped into a CI type (CIT). Each CIT provides a template for creating a new CI and its associated properties.

A relationship defines the link between two CIs. Relationships represent the dependencies and connections between the entities in your IT environment. As an example, in the diagram below all CITs that can connect to CIT host are listed with the permitted relationships between them.
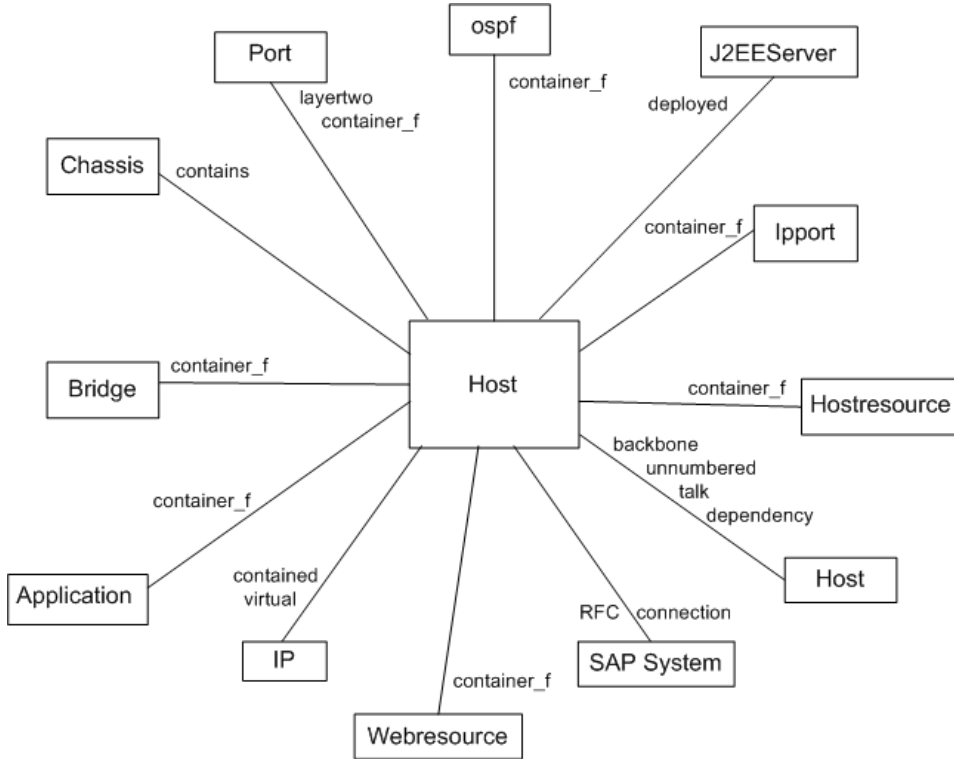
Diagram of CITs Whose Parent is System

The following diagram depicts all CITs whose immediate parent is system. To view the entire list of CITs under system, see "List of Parent, Description, and Key Attributes of all CI Types," below.
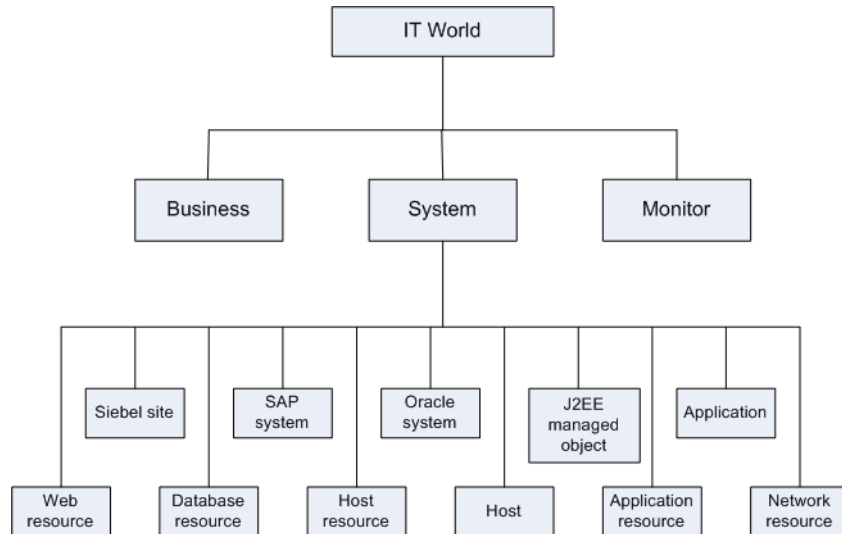
List of Parent, Description, and Key Attributes of all CI Types

The following table lists all CITs defined in the CIT Manager, the CIT's parent, a brief description, and key attributes if they exist.

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
agent	application	Virtual class to all agents.	
apache	webserver	Apache Web Server.	webserver_configfile

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
apachevirtualhost	applicationresource	Apache Virtual Host.	
oracleapplication	oracleappsresource	Oracle E-Business Suite Application.	
applicationresource	system	Virtual class to all application resources.	root_container, data_name
atmport	networkresource	A port on an Asyschronous Transmit Mode switch.	atmport_pnniport number, atmswitching_pn nninodeaddress
atmswitch	host	Asynchronous Transfer Mode switch:a network technology that supports realtime voice, video, and data. The topology uses switches that establish a logical end-to-end circuit.	
atmswitching	networkresource	Asynchronous Transfer Mode switching.	ress
bridge	networkresource	A device that connects one local area network (LAN) to another.	bridge_basemacad dr
business	it_world	Business classes.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
chassis	networkresource	A chassis that holds Asynchronous Transfer Mode (ATM) switches.	chassis_uniqueid
ipclient	ipport	The client port to which the server port is connected.	ipclient_remoteip, ipclient_remoteport
concentrator	host	A switch or a hub with no SNMP agent.	host_key, data_name
configfile	document	Holds data from a configuration file of an application.	
connectionpool	jdbcresource	A cache of database connections maintained in the database's memory so that the connections can be reused when the database receives future requests for data. Connection pools are used to enhance the performance of executing commands on a database.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
cpu	hostresource	Central Processing Unit.	cpu_cid
csqlprocesses	sqlprocesses	An CI representing a Complete SQL process.	
database	application	A collection of data that is organized so that its contents can be easily accessed, managed, and updated.	
databaseresource	system	A virtual class to all database resources.	root_container, data_name
dbarchivefile	file	An archive file for a Microsoft SQL Server 2000 Analysis Services database. It contains the contents of the directory whose name is the same as the database name.	
dbcontrolfile	file	A database file containing information needed to maintain and verify database integrity.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
dbdatafile	databaseresource	The database file that contains the database's data.	dbdatafile_fileid
dbextent	databaseresource	A additional unit of space allocated to a SQL Server CI, such as a table or index, whenever the CI needs more space.	dbextent_segment type, dbextent_segment name, dbextent_extentid, dbextent_owner
dbindex	databaseresource	A mechanism to locate and access data within a database. An index may refer to one or more columns and be a means of enforcing uniqueness in column values.	dbindex_owner
dbjob	databaseresource	A specified series of operations performed sequentially by the database manager.	dbjob_jobid

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
dblinkobj	databaseresource	The link between databases that stores information about how to connect to the remote database.	
dbredofile	file	A transaction log that records data modifications made in the database.	
dbsegment	databaseresource	Contains the component database and any utilities provided by the developers for the DBA's use in installing and filling that particular database.	dbsegment_owner, dbsegment_segmenttype
dbsnapshot	databaseresource	A backup that can be created from an entire database or individual files.	
dbtable	databaseresource	Organizes the information about a single topic into rows and columns.	dbtable_owner

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
dbtablespace	databaseresource	Represents the amount of space required to store the data in a table.	
dbuser	databaseresource	A defined user in the database.	
db2	database	An IBM database management system that serves a number of different operating system platforms.	
dbaobjects	databaseresource	An CI that represents a database procedure, function, package, and package-body.	dbaobjects_type
daemon	hostresource	A program that forwards requests to other programs or processes.	
dir	file	The name of a group of files in the same location.	file_path
disk	hostresource	A logical division of a hard disk.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
document	applicationresource	A physical file.	
drive	logicaldisk	A hard disk drive.	
ejb	j2eemanagedobject	EJB component in Sun's J2EE platform that provides a Java environment for developing and running distributed applications.	
ejbmodule	j2eemodule	A module containing Enterprise JavaBeans.	
elan	networkresource	Emulated Local Area Network:can be created using ATM LAN emulation technology to join together existing LANs running protocols such as IP, Novell IPX, AppleTalk, and DECnet.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
enterasyse7blade	host	The atmswitch class represents a network element which its oid represents an enterasyse7blade.	
entitybean	ejb	An Enterprise JavaBean entity.	
eum_monitor	monitor	Monitor system classes.	
eventlog	hostresource	Microsoft event log data.	eventlog_type, eventlog_application
file	hostresource	An entity of data.	
filesystem	logicaldisk	UNIX file system.	
firewall	host	A gateway that limits access between networks in accordance with local security policy.	
host	system	A network element that has a unique IP address.	host_key
hostresource	system	Virtual class to all host resources.	root_container,data_name

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
httpcontext	applicationresource	An IBM HTTP server plug-in.	root_container, data_name
ibmhttpserver	webserver	An IBM HTTP server.	webserver_configfile
mqaliasq	mqqueuelocal	The Message Queue alias queue to a local queue or to a remote queue definition.	root_container, data_name
mqchannel	applicationresource	Message Queue channel data.	root_container, data_name
mqchclntconn	mqmqichannel	Message Queue client connection channel for Message Queue Interface (MQI) calls.	root_container, data_name
mqcluster	applicationresource	Message Queue cluster data.	
mqchclusrcvr	mqmsgreceiverchannel	Message Queue cluster receiver channel.	root_container, data_name
mqchclussdr	mqmsgsenderchannel	Message Queue cluster sender channel.	root_container, data_name
mqmsgchannel	mqchannel	Message Queue channel data.	root_container, data_name
mqmsgreceiverchannel	mqmsgchannel	Message Queue channel data.	root_container, data_name
mqmsgsenderchannel	mqmsgchannel	Message Queue channel data.	root_container, data_name

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
mqmqichannel	mqchannel	Message Queue channel data.	root_container, data_name
mqqueue	applicationresource	Message Queue definition data.	
mqqueuelocal	mqqueue	A local Message Queue queue.	root_container, data_name
mqqueuemanager	applicationresource	Message Queue manager data.	
mqqueueremote	mqqueue	A remote Message Queue queue definition.	root_container, data_name
mqchrcvr	mqmsgreceivercha nnel	Message Queue receiver channel.	root_container, data_name
mqchrqstr	mqmsgreceivercha nnel	Message Queue requester channel.	root_container, data_name
mqchsdr	mqmsgsenderchan nel	Message Queue sender channel.	root_container, data_name
mqchsvr	mqmsgsenderchan nel	Message Queue server channel.	root_container, data_name
mqchsvrconn	mqmqichannel	Message Queue server connection channel for Message Queue Interface (MQI) calls.	root_container, data_name
mqxmitq	mqqueuelocal	An Message Queue transmission queue.	root_container, data_name

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
webspheremq	application	Message Queue for WebSphere data.	
iis	webserver	Internet Information Server.	
interface	networkresource	A physical and logical way to support an attachment of any device to a connector or to another device.	interface_macaddr
interfaceindex	networkresource	Represents the unique number of an interface.	interfaceindex_index
ip	networkresource	Internet Protocol:the address which identifies the sender or the receiver of information sent across the internet.	ip_address, ip_domain
ipfw	ip	The IP address of a firewall.	ipfw_hopnumber, ipfw_trailid
ipunknown	ipport	Unknown TCP/UDP port.	
j2eeapplication	j2eedeployedobject	An application inside the application server	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
j2eecluster	j2eeresource	A group of application servers that run a J2EE application as if they were a single server.	data_name
j2eedeployedobject	j2eemanagedobject	J2EE deployed CI.	root_container, data_name
j2eedomain	j2eemanagedobject	An application server domain.	
executequeue	j2eeresource	executequeue	data_name
j2eemanagedobject	system	A virtual class to all J2EE resources.	root_container, data_name
j2eemodule	j2eedeployedobject	J2EE deployed CI.	
j2eeresource	j2eemanagedobject	A virtual class to all J2EE resources.	
j2eeserver	application	A server within a J2EE application server.	
jdbcdatasource	j2eemanagedobject	Java Database Connectivity data source.	
jdbcprovider	jdbcresource	A provider for Java Database Connectivity.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
jdbcresource	j2eeresource	A virtual class to all Java Database Connectivity resources.	
jmsdatastore	jmsresource	Java Messaging Service data store.	
jmsdestination	jmsresource	Java Messaging Service destination.	
jmsprovider	jmsresource	A provider for Java Messaging Service.	
jmsresource	j2eeresource	A virtual class to all Java Messaging Service resources.	
jmsserver	jmsresource	Java Messaging Service server.	
JVM	j2eemanagedobject	Java Virtual Machine.	
lb	host	Load Balancer:distributes processing and communication activity evenly across a network so that no single device is overwhelmed.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
logdir	dir	Log directory definition.	file_path
logfile	file	Requested information which is kept in a file.	
logicaldisk	hostresource	A logical disk.	
lpar	host	Logical partition: the division of a computer's processors, memory, and storage into multiple sets of resources so that each set of resources can be operated independently with its own operating system instance and applications.	
mainframe	host	Main computer that can serve many users.	
marconiatmswitch	atmswitch	AMarconi ATM switch.	
memory	hostresource	Random access memory (RAM).	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
messagedrivenbean	ejb	Enterprise JavaBeans in Sun's J2EE platform. Message driven beans are generated to process Java Messaging Service (JMS) messages.	
mibtree	hostresource	Management Information Base tree:a database of CIs that can be monitored by a management system. SNMP is a management protocol of MIB.	mibtree_oid
monitor	it_world	Monitoring classes.	
msdomain	networkresource	Represents Microsoft domains and workgroups in a Microsoft network.	
netdevice	host	A network device.	
netprinter	netdevice	A printer device that serves the network's users.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
network	networkresource	A series of elements or nodes connected to each other by communication paths.	network_domain, network_netmask, network_netaddr
networkresource	system	A virtual class to all network resources.	root_container, data_name
nt	host	Host with Microsoft operation system NT.	
service	hostresource	Software that can be controlled by the Microsoft Management Console (MMC).	
ntcmd	shell	xcmd or pstools. NT command line tools:allows you to execute applications on remote systems without installing any client software. Examples of such tools are xCmd and PsTools.	
oracle	database	Oracle database.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
oraclesystem	system	Oracle E-Business Suite.	oraclesystem_dba ddress
oracleappsresource	applicationresource	Oracle E-Business Suite Resource.	
oracleias	j2eeserver	Oracle Internet Application Server.	
osuser	hostresource	Operating system user.	
ospf	networkresource	Open Shortest Path First (OSPF) protocol:an interior gateway protocol used to distribute routing information within a single autonomous system.	ospf_areaid
port	networkresource	A place which being physically connected to another device.	port_number, root_container
printq	hostresource	A printer queue.	
printer	hostresource	A printer device.	
process	hostresource	An instance of a program.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
program	process	Name of an installed application or software.	data_name
ras	host	RAS (row address strobe) is a signal which being sent to a dynamic random access memory (DRAM Row Address Strobe:a clock signal in a dynamic random access memory (DRAM) chip used to pinpoint the row of a particular bit in a row-column matrix Remote Access Server	
router	host	A device or software connected to at least two networks. It decides the next network point to which to send the information packet.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
sap_application_component	sap_resource	A group of SAP transactions with common application components.	type
sap_app_server	application	SAP Application Server.	
sap_business_process	sap_resource	SAP Business Process..	
sap_business_scenario	sap_resource	SAP Business Scenario.	
sap_client	sap_resource	SAP Client.	
sap_gateway	application	SAP Gateway.	
sap_its_agate	application	SAP Internet Transaction Server AGate:an application gateway that establishes a link to the R/3 system and performs the processing tasks required to move data between an R/3 application and the Internet.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
sap_its_wgate	sap_resource	SAP Internet Transaction Server WGate:a web gatewaythat establishes a connection between Internet Transaction Server (ITS) and the web server. It forwards user requests to the AGate.	
sap_j2ee_app_server	sap_app_server	Represents an instance of the J2EE Application server.	
sap_j2ee_dispatcher	sap_resource	Receives a client request and forwards it to the server process with the lowest capacity usage.	
sap_j2ee_server_process	sap_resource	A single J2EE server process instance.	
sap_process_step	sap_resource	SAP Process Step.	
sap_bp_project	sap_resource	SAP Business Blueprint Project.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
sap_r3_server	sap_app_server	SAP's integrated software solution for client/server and distributed open systems.	
sap_resource	applicationresource	SAP Resource.	
sap_system	system	A logical unit that groups SAP-related entities and possibly other entities into one homogenous SAP deployment.	data_name
sap_transaction	sap_resource	A business process defined in the SAP System.	root_container, data_name
sap_transport	sap_resource	SAP Transport.	
sap_transport_cha nge	sap_resource	SAP Transport Change.	
sap_work_process	sap_resource	A logical single-instance representation of a certain type of work process.	
ipserver	ipport	A known server port listening to or connected to a client port.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
oracleappservice	oracleappsresource	Oracle E-Business Suite Application Service.	
oracleappservicemanager	oracleappsresource	Oracle E-Business Suite Service Manager.	
servlet	j2eemanagedobject	A server-side Java program that provides additional features to the server.	
sessionbean	ejb	An Enterprise JavaBean session.	
shell	application	xcmd or pstools	
siebel_app_server	application	A server application running the Siebel business logic tier.	root_container, data_name
siebel_component	applicationresource	A process on the Siebel Application Server incorporating some Siebel application functionality.	alias

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
siebel_comp_grp	applicationresource	An collection of Siebel components running on the Siebel application server.	alias
siebel_gateway	application	A coordinating server that routes requests to the correct component.	
siebel_site	system	A logical unit describing a group of servers functioning together to build a complete Siebel toolset experience.	gateway_address, data_name
siebel_web_app	applicationresource	A Siebel application's location on the web server.	
siebel_wse	applicationresource	Web Server Extension:a supported add-on to a web server.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
snmp	agent	Simple Network Management Protocol:protocol used by network hosts to exchange information used in network management.	snmp_port
software	hostresource	The general name of a program.	data_name
application	system	A virtual class to all applications.	root_container, data_name
sqlalert	databaseresource	SQL Server Agent compares events to user-defined alerts. When the Agent finds a match, it fires an alert.	sqlalert_alertid
sqlbackup	databaseresource	A backup file generated by the system or by user requests.	
sqldatabase	databaseresource	A SQL database from Microsoft.	
sqlfile	databaseresource	A SQL file.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
sqljob	databaseresource	A specified series of operations performed sequentially by a SQL Server Agent. A job can perform a wide range of activities, including running Transact-SQL script, command line applications, and Microsoft ActiveX scripts.	sqljob_jobid
sqljobstep	databaseresource	An action that the job performs on a database or a server. Every job must have at least one job step.	sqljobstep_stepid
sqlperformancemonitor	databaseresource	Performance monitor constants that describe Microsoft Windows NT Performance Monitor polling behavior.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
sqlprocesses	applicationresource	SQL processes.	sqlprocesses_host name, sqlprocesses_data basename, sqlprocesses_progr am
sqlserver	database	SQL server.	
ssh	shell	Secure Shell:an agent that provides secure encrypted communication between two untrusted hosts over an insecure network.	
statefulsessionbean	sessionbean	EJB stateful:an Enterprise JavaBean that is associated with one client at a time. It is assigned to a new client only after a previous client releases it or it times out.	
statelesssessionbea n	sessionbean	EJB stateless:an Enterprise JavaBean that is context independent. It does not maintain a state between session calls.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
sunoneserver	webserver	Sun ONE Web Server.	
switch	host	A device that sends incoming data from any of multiple input ports to a specific output port in order to move the data toward its intended destination.	
switchrouter	switch	A switch that also acts as a router.	
sybase	database	Sybase.	
sybasedb	databaseresource	Sybase database.	
sysplex	networkresource	SYStem comPLEX:the multiprocessing capability of IBM MVS/ESA and OS/390 mainframes.	
system	it_world	System classes.	
system_monitor	monitor	Monitor system classes.	
ipport	networkresource	A logical connection between ports using TCP/IP Internet Protocol.	ipport_number

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
telnet	shell	A telnet connection for TCP/IP networks that allows you to execute programs on a remote computer.	
terminalserver	host	A device that provides terminals with a common connection using RS232 protocol.	
tomcat	webserver	A web server from the Apache Software Foundation that executes Java servlets and displays web pages that have embedded Java code.	
trail	networkresource	The Trail element describes the path from one network to another.	trail_destnetmask, trail_destnetaddr, trail_srcnetdomain, trail_srcnetaddr, trail_srcnetmask, trail_destnetdomain
unix	host	UNIX operating system.	

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
url	webresource	Uniform Resource Locator:the address of a resource on the Internet.	url_connectstring
vax	host	VAX:Virtual Address eXtension machine.	
vlan	networkresource	Virtual LAN:a network of computers that behave as if they are connected to the same wire even though they may be physically located on different segments of a LAN.	vlan_number
webapplication	j2eedeployedobject	Web application.	
oraclewebcomponent	oracleappsresource	Oracle E-Business Suite Web Component.	
webmodule	j2eemodule	J2EE deployed CI.	
webresource	system	A virtual class to all web resources.	root_container, data_name

CI TYPE NAME	PARENT	DESCRIPTION	KEY ATTRIBUTES
webserver	application	A web server.	
websphere	application	A WebSphere application server based on a Java 2 platform.	root_container, data_name
wmi	agent	Windows Management Instrumentation :a connection allowing scripts to monitor and control managed resources throughout the network.	
xterminal	host	A device using X11 emulation to connect to a UNIX machine.	
db2user	databaseresource	A user in a DB2 database.	

Permitted Relationships between CI Types

The following table lists all CI types defined under CIT system and their permitted relationship with other CI types. For more information about CI types, see "List of Parent, Description, and Key Attributes of all CI Types" on page 501. The definition of each relationship can be found in Relationship Definitions.

CI Type 1	CI Type 2	Relationship
application	applicationresource	container_f
atmport	atmport	pnniconnection
	interfaceindex	parent
	ip	uniconnection
atmswitch	atmswitching	container_f
atmswitching	atmport	container_f
bridge	port	container_f
	port	contains
chassis	elan	chassiselanmap
concentrator	port	layertwo
connectionpool	database	depend

CI Type 1	CI Type 2	Relationship
database	daemon	depend
	dbaobjects	container_f
	dbarchivefile	container_f
	dbcontrolfile	container_f
	dbdatafile	container_f
	dbjob	container_f
	dblinkobj	container_f
	dbredofile	container_f
	dbsegment	container_f
	dbsnapshot	container_f
	dbtablespace	container_f
	dbuser	container_f
	process	dbclient
	service	depend
dbdatafile	dbextent	container_f
dbindex	dbextent	resource
dblinkobj	database	dblink
dbsegment	dbextent	resource
dbsnapshot	dbjob	depend
	dblinkobj	resource
dbtable	dbextent	resource
dbtablespace	dbdatafile	resource
	dbindex	container_f
	dbsegment	resource
	dbtable	container_f

CI Type 1	CI Type 2	Relationship
dbuser	dbaobjects	owner
	dbextent	owner
	dbindex	owner
	dbjob	owner
	dblinkobj	owner
	dbsnapshot	owner
	dbtable	owner
dir	file	container_f
drive	drive	share
ejbmodule	ejb	container_f
elan	chassis	chassisvlanmap
	switch	bcastdomain
filesystem	filesystem	nfs

CI Type 1	CI Type 2	Relationship
host	application	container_f
	bridge	container_f
	chassis	contains
	host	backbone
	host	unnumbered
	host	talk
	host	dependency
	hostresource	container_f
	interface	container_f
	interfaceindex	container_f
	ip	virtual
	ip	contained
	ipport	container_f
	j2eeserver	deployed
	ospf	container_f
	port	layertwo
	port	container_f
	webresource	container_f
httpcontext	j2eeapplication	use
interface	interface	usb
	port	layertwo
interfaceindex	interface	parent

CI Type 1	CI Type 2	Relationship
ip	interface	parent
	interfaceindex	parent
	ip	depend
	ip	route
	ip	traffic
	ipserver	clientserver
ipport	ip	use
	ipport	tcp
	ipserver	clientserver
j2eeapplication	j2eemodule	container_f
	jdbcdatasource	use
j2eecluster	database	member
	j2eemanagedobject	container_f
	j2eeserver	member
	sap_j2ee_app_server	member
j2eedomain	j2eemanagedobject	container_f
	j2eeserver	member
j2eemanagedobject	jvm	use

CI Type 1	CI Type 2	Relationship
j2eeserver	executequeue	use
	ip	j2eesocket
	ip	depend
	ipport	depend
	j2eemanagedobject	deployed
	j2eemanagedobject	container_f
	j2eeserver	container_f
	jvm	use
jdbcdatasource	connectionpool	depend
	database	depend
jdbcprovider	jdbcdatasource	container_f
jmsserver	jmsdatastore	container_f
	jmsdestination	container_f
logdir	logfile	container_f
	logfile	parent
logicaldisk	dir	container_f
	disk	parent
	file	container_f
mainframe	lpar	member
	sysplex	member
mqchannel	mqchannel	mqmsglink
mqchclusrcvr	mqchclussdr	mqmsglink
mqchrcvr	mqchsdr	mqmsglink
	mqchsvr	mqmsglink

CI Type 1	CI Type 2	Relationship
mqchrqstr	mqchsdr	mqmsglink
	mqchsvr	mqmsglink
mqchsvrconn	mqchclntconn	mqmqilink
mqcluster	mqchannel	member
	mqqueue	member
	mqqueuemanager	mqrepository
	mqqueuemanager	member
mqmsgreceiverchannel	mqmsgsenderchannel	mqmsglink
mqqueue	mqaliasq	mqalias
	mqqueue	mqresolve
mqqueuelocal	mqqueueremote	mqresolve
mqqueuemanager	mqchannel	container_f
	mqqueue	container_f
mqqueueremote	mqxmitq	use
mqxmitq	mqchannel	mqchannelof
msdomain	host	member
network	host	member
	ip	member
nt	drive	container_f
	service	container_f
oracle	dbtns	container_f
oracleapplication	dbtablespace	use
	oracleapplication	depend
oracleappservice	process	resource
oracleappservicemanager	oracleappservice	member

CI Type 1	CI Type 2	Relationship
oracleias	oracleappsresource	deployed
	oracleappsresource	container_f
oraclesystem	oracle	member
	oracleappsresource	container_f
	oracleias	member
osuser	dir	container_f
port	bridge	contained
	bridge	bridgebackbone
	interfaceindex	parent
	port	backbone
	vlan	vlanmembership
printer	printq	parent
printq	printq	parent
process	daemon	brother
	file	use
	ipport	use
	printq	use
	service	brother
sap_application_component	sap_application_component	container_f
	sap_transaction	contains
sap_bp_project	sap_business_scenario	container_f
sap_business_process	sap_process_step	container_f
sap_business_scenario	sap_business_process	container_f
sap_its_agate	sap_its_wgate	depend

CI Type 1	CI Type 2	Relationship
sap_j2ee_server_process	j2eemanagedobject	deployed
sap_process_step	sap_transaction	execute
sap_r3_server	configfile	use
	sap_its_agate	depend
	sap_work_process	container_f
sap_system	database	depend
	host	sap_rfc_connection
	sap_gateway	member
	sap_r3_server	member
	sap_resource	container_f
sap_transport	sap_transaction	use
	sap_transport_change	container_f
sap_transport_change	sap_transaction	use
service	ipport	use
	service	depend
	sqlserver	depend
siebel_app_server	database	depend
	siebel_comp_grp	container_f
	siebel_gateway	depend
siebel_comp_grp	siebel_component	container_f
siebel_component	siebel_web_app	depend
siebel_gateway	siebel_wse	depend
siebel_site	host	contains
	siebel_app_server	member
	siebel_gateway	member

CI Type 1	CI Type 2	Relationship
siebel_wse	siebel_web_app	container_f
snmp	mibtree	container_f
sqldatabase	sqlalert	deployed
	sqlbackup	container_f
	sqlfile	container_f
	sqlperformancemonitor	container_f
	sqlprocesses	deployed
sqlfile	disk	depend
sqljob	sqljobstep	container_f
sqlprocesses	process	dbclient
sqlserver	sqlalert	container_f
	sqldatabase	container_f
	sqljob	container_f
	sqlperformancemonitor	container_f
	sqlprocesses	container_f
switch	vlan	container_f
sybase	sybasedb	container_f
sybasedb	process	dbclient
sysplex	lpar	member
unix	daemon	container_f
	filesystem	container_f
vlan	bridge	vlantobridge
	elan	elanvlanmap
webapplication	servlet	container_f

CI Type 1	CI Type 2	Relationship
webmodule	servlet	container_f
webserver	sap_its_wgate	container_f
	siebel_wse	container_f
websphere	j2eecluster	container_f
	j2eeserver	container_f
webspheremq	mqcluster	container_f
	mqqueuemanager	container_f

B

Relationship Definitions

This document defines relationships used throughout Mercury Application Mapping documentation.

backbone

Represents a physical connection between two switches. The relationship is discovered by the Discover layer 2 service.

bcastdomain

The relationship between an elan (emulated LAN) and a switch.

bridgebackbone

Represents a physical connection between two switches connecting a switch port to a switch bridge. The relationship is discovered by the Discover base service.

brother

The relationship among elements who share the same parent.

chassiselanmap

The relationship between a chassis and an elan (emulated LAN).

chassisvlanmap

The relationship between a chassis and a vlan (virtual LAN).

clientserver

Represents a row of data from the tcpConnLocalAddress table in the Management Information Base (MIB) tree. This data has information about the TCP connection between two hosts' ports when you can differentiate between the server port and the client port. The tcpConnLocalAddress table is in the MIB address 1.3.6.1.2.1.6.13.1.2. The clientserver relationship is discovered by the Discover TCP connection collector.

contained

The relationship between two CIs whereby a second CI is included in the first CI. This relationship is found only between IP and host.

container_f

The functional relationship between a parent and a child. The child does not inherit any properties.

contains

The relationship between two CIs whereby a second CI is included in the first CI.

dbclient

The relationship between a process and a database.

dblink

The relationship between a database and a database link object.

depend

The relationship wherein one CI needs a functionality of another CI.

dependency

The relationship wherein one CI needs a functionality of another CI.

deployed

The relationship wherein one CI is put into action by another CI.

elanvlanmap

The relationship between elan (emulated LAN) and vlan (virtual LAN) network components.

execute

> The relationship between an agent and the job it deploys.

j2eesocket

> The relationship between a server and a remote client.

layertwo

> Represents the physical connection between a switch and a host. It is discovered by the Discover layer 2 service.

member

> The relationship between two CIs whereby one CI is included in another CI.

mqalias

> The relationship between an alias queue and a local queue.

mqchannelof

> The relationship between a channel and its transmission queue.

mqmqilink

> The relationship between client and server channels for Message Queue Interface (MQI) calls.

mqmsglink

> The relationship between two channels for message transfer.

mqrepository

> The relationship between a message queue cluster and a message queue queue manager.

mqresolve

> The relationship between a remote queue and the local queue to which it is mapped.

nfs

> Network file server.

owner

> The user of a resource.

parent

The relationship between elements where one element is parent of the other. For example, **ip** is the parent of **interface**.

pnniconnection

The relationship between two ATM ports.

resource

The relationship between elements where one element is the resource of the other. For example, **dbsnapshot** is the resource of **dblinkobject**.

route

Represents a row of data from the routing table in the Management Information Base (MIB) containing the data of the next_hop ip address and the destination network address. The routing table is in the MIB address 1.3.6.1.2.1.4.21.1.7. The relationship is discovered by the Discover route collector.

sap_rfc_connection

The relationship between a SAP system and a host. The host may be another SAP system or a non-SAP system.

share

The relationship between two drives.

talk

The relationship between two hosts.

tcp

Represents a row of data from the tcpConnLocalAddress table in the Management Information Base (MIB). This data has information about the TCP connection between two hosts' ports when you cannot differentiate between the server port and the client port. The tcpConnLocalAddress table is in the MIB address 1.3.6.1.2.1.6.13.1.2. The relationship is discovered by the Discover TCP connection collector.

traffic

Represents all network flow, regardless of protocol, between two IPs.

uniconnection

The relationship between an ATM port and an IP.

unnumbered

Represents a row of data from the routing table in the Management Information Base (MIB). This data has information about the next_hop IP address and the destination network address. The routing table is in the MIB address 1.3.6.1.2.1.4.21.1.7. Unnumbered relationships are discovered by the base collector.

usb

The relationship between two interfaces.

use

The relationship between elements whereby one element uses the other. For example, **process** uses **file**.

virtual

The relationship between a router and its virtual IP.

vlanmembership

The relationship between a vlan (virtual LAN) and a physical port.

vlantobridge

The relationship between a vlan (virtual LAN) and a bridge.

Appendixes

C

Glossary

This document lists terminology that is used throughout the Mercury Application Mapping documentation.

acknowledge

> Indicates that an event that has occurred in the system has been recognized. The acknowledgment action is manifested through the stopping of the blink of the CI that is related to the event.

active event

> An event that is displayed on the Event Browsers and Tabs, and manifested via state changes and icons blinks in the Topology Map. It is created when a raw event meets a certain condition, and as a result is converted to an active event.

admin state

> The administrative state of a selected CI (Disabled, Restricted, Testing, Unknown, Managed, or Unmanaged). Any value other than **Managed** indicates that any event information coming from the selected CI is not reliable because the CI is disabled, undergoing testing or is otherwise unmanaged. The admin state of a CI does not propagate upwards to CIs at higher levels.

attribute

> A specific characteristic of a selected CI that is defined in the configuration item type model and belongs to one of following categories: Key, Attribute, Dynamic Attribute.

auto-discovery

A type of discovery where a discovery method automatically detects resources that are not previously known.

availability rules

Rules that are defined for views and nodes, which enable calculating and monitoring the availability of managed CIs during predefined periods. The availability of a CI is measured as a percentage of "up time" versus "total time."

CI

An entity in a managed world that is distinguished from other CIs by its CIT and key attributes.

CIT

An entity contained in the configuration item type model.

configuration item type model

The set of defined topology CIs and their relationships that describes a managed world.

compound

The children of a selected CI.

compound status

Exists in parent CIs only and indicates the maximum status of the selected CI's children. This status does not include the map status of the parent CI.

container relationship

A functional relationship that connects a host and a CI that is fully dependent on the host.

correlation rule

A rule that specifies a casual relationship between two nodes.

correlation state

> A state that is defined in a correlation rule and assigned to an affected CI once its root-cause CI fulfills the correlation condition and a correlation event occurs. There are two ways to define this state: 1- It can be fixed in advance (for example. when the correlation event occurs, the Corr State of the affected CI is always Critical); 2- It can be defined as percentage of the Oper State of its root-cause CI (for example, when the correlation event occurs, the Corr State of the affected CI is 50% of the root-cause CI's Oper State.) When several correlation events are simultaneously related to a CI, the correlation state is the maximum severity of all the events.

discovery

> The process of finding resources within an IT infrastructure.

Discovery Management

> A subsystem, whose function is to manage the discovery process by defining the discovery scope, defining the protocol, defining the connection data for the protocol, and activating the discovery pattern(s).

discovery method

> The component that performs the actual discovery. Each discovery method is designed to discover a specific resource. The methods can be used individually or combined with each other to collect data from and about IT resources.

discovery pattern

> An XML template which contains invocation definitions of one or more discovery tasks. The template defines which data is collected, through which protocol, when the data should be passed to the Mercury Universal CMDB, and the executing schedule of the pattern.

Discovery Probe

> A component that hosts discovery methods and provides management services that enable their manipulation. The Discovery Probe receives a task, delivers its parameters to a discovery method, manages the discovery method performance, and returns the task results to the Discovery Manager.

discovery task

The executable part of a discovery pattern, which triggers discovery methods to begin discovering the data defined in the pattern. Several discovery tasks can be defined in one pattern and executed from it. The discovery tasks include static parameters that are directly received from the discovery pattern (such as threshold and schedule), as well as runtime parameters that are obtained from the data that is collected by the pattern's trigger.

derived (node)

A view's node definition that determines that both the selected CI and any CIs derived from the selected CI (inheritance CIs) are displayed in the topology map.

View explorer

Displays a hierarchical tree structure of defined items (queries, views, and so forth).

event

Some change that occurs in the managed world.

event filter

A predefined condition that filters events displayed in the Event Browser, Tabs and Log. When applied only those events that meet the criteria defined by the selected filter are displayed.

event message

A predefined message sent in response to the occurrence of particular events.

event CI

A CI created from a raw event that contains all the attributes of the event, including the CI to which it is related.

event severity

An attribute of an event CI that indicates the severity of the event (Normal, Warning, Minor, Major, Critical). Each severity level is displayed in a different color. This event severity determines the operational state of the related CI.

filter editor

A tool that enables the creation and use of existing event filters to reduce the number of events displayed in the Event tabs, browsers and log.

inheritance relationship

Classes that inherit the attributes and behavior (such as the label function) of the CIT above it. For example, the IPserver and IPclient classes inherit attributes from the IPport CIT.

interdependencies

See the glossary item **relationship**.

join_f link

A relationship that connects two CIs that have at least one comparable value. For example, it is possible to compare the values of IP domain and network domain, and host DNS and IP address.

label

The title that appears under a CI in a map. Label definition can be customized to include different attribute values. For example, if in host the function label is composed of hostname & network, the displayed label will be: server1 10.0.65.0.

layer

Each map in MAM can be represented in a layered structure. The CIs are divided into layers and you can drill down from one layer to the next to see more CIs.

relationship

Also known as interdependencies. An entity that defines the relationship between two CIs. Relationships can describe either a physical or logical connection between two CIs, or a functional relationship.

managed world

Any self-contained environment that can be described using a topology model.

managed view

A series of rules and definitions for displaying query results in Mercury Application Mapping.

management categories

Categories, such as performance and change, with which you can manage any business service.

Map report

Report that process and display data of views and nodes, according to the way they are arranged and expressed in the Topology Map context. The Map report can produce and display statistics about the levels of a view or a node in one report.

map status

The operational state of the CI, as adjusted for the weight accorded to its significance in the system. (The significance weight is set as the node's status factor.) For example, if the state (the maximum severity) of the event is **5** and the status weight is defined at 100%, the CI's map status is **5** (5*100/100). If the weight is defined as 60%, the CI's status is **3** (5*60/100). If the CI has children, the map status value is the maximum status of the CI and all its children.

oper state

The current operational state of the selected CI, which is the maximum severity of any event currently occurring to the CI.

organizational (folding) rules

Rules that define how linked CIs are displayed in the topology map. There are four organizational rules: parent, child, right sibling and left sibling.

patterns

See discovery pattern.

priority indicators

> Different priority levels of views, which are expressed in the colors of the view icons: high = yellow, medium = purple, low = grey. Priority levels are defined as part of a view's properties, and can help differentiate between the types of attention and support that each view requires.

propagation

> The transmission to higher view levels of managed CI data, such as status or blink.

raw event

> A message that notifies Mercury Application Mapping that a specific change has occurred in a managed CI.

resources

> Applications, databases, network devices (hardware and software), and servers that are discovered and managed by Mercury Application Mapping.

severity

> Indicates the gravity level of an event (Normal, Warning, Minor, Major, Critical). Each severity level is displayed in a different color, and it determines the Operational State of the CI that is related to the event.

status factor

> A value (in percent) that defines the significance of the CI status. This is useful for reducing the significance of CIs whose status can fluctuate on a regular basis, such as processes, which continuously change. By reducing the significance of such CIs in the system, events are prevented from generating an overflow of critical messages. For example, if the state of an event is **5** and the status weight of the process is defined at 40%, then the CI status is **2** (5*40/100).

system report

> A report that is based on a specific Report TQL query. The report display statistics about a query results and can be defined to display very detailed and focused information.

tags

> Logical groupings of CIs and relationships that can be created without any prerequisites, that is, the grouped CIs do not need any prior common denominator. The tags are additional organizing filters for presentation purposes.

TQL

> A language and a tool for discovering, organizing, and managing IT infrastructure data. TQL is an abbreviation for Topology Query Language. It a language that extends standard SQL, by adding the ability to draw conceptual relationships between CIs and create visual images of IT infrastructure resources.

topology

> A visual description of the configuration or arrangement of an IT infrastructure, including its nodes and connecting relationships.

Mercury Universal CMDB

> The core information repository of Mercury Application Mapping. It stores and handles the infrastructure data collected and updated by the discovery system. The information concerning discovered CIs and relationships is deposited, grouped, and updated in the form of CIT definitions according to CI-oriented methodology.

topology CI

> Any member of a CIT that exists within the topology of the managed world.

view

> A collection of CIs and relationships represented by icons. These CIs and relationships are the result of a TQL query, and are displayed as a view according to display and organizational rules that are assigned to them. Each CI/relationship can be presented in multiple views or by multiple icons in different view layers.

viewing layer

> A two-dimensional arrangement of topology CIs to be viewed by a user.

Index