MERCURY
APPLICATION MAPPING™

User's Guide

MERCURY™

Mercury Application Mapping
User's Guide
Version 3.0

MERCURY™

Mercury Application Mapping User's Guide, Version 3.0

This manual, and the accompanying software and other documentation, is protected by U.S. and international copyright laws, and may be used only in accordance with the accompanying license agreement. Features of the software, and of other products and services of Mercury Interactive Corporation, may be covered by one or more of the following patents: United States: 5,511,185; 5,657,438; 5,701,139; 5,870,559; 5,958,008; 5,974,572; 6,137,782; 6,138,157; 6,144,962; 6,205,122; 6,237,006; 6,341,310; 6,360,332, 6,449,739; 6,470,383; 6,477,483; 6,549,944; 6,560,564; 6,564,342; 6,587,969; 6,631,408; 6,631,411; 6,633,912; 6,694,288; 6,738,813; 6,738,933; 6,754,701; 6,792,460 and 6,810,494. Australia: 763468 and 762554. Other patents pending. All rights reserved.

Mercury, Mercury Interactive, the Mercury logo, the Mercury Interactive logo, LoadRunner, WinRunner, SiteScope and TestDirector are trademarks of Mercury Interactive Corporation and may be registered in certain jurisdictions. The absence of a trademark from this list does not constitute a waiver of Mercury's intellectual property rights concerning that trademark.

All other company, brand and product names may be trademarks or registered trademarks of their respective holders. Mercury disclaims any responsibility for specifying which marks are owned by which companies or which organizations.

Mercury Interactive Corporation
379 North Whisman Road
Mountain View, CA 94043
Tel: (650) 603-5200
Toll Free: (800) TEST-911
Customer Support: (877) TEST-HLP
Fax: (650) 603-5300

© 2004-2005 Mercury Interactive Corporation, All rights reserved

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

MAMUG3.0/01

Table of Contents

PART III: VIEWS, ORGANIZATION RULES, AND NOTIFICATIONS

PART IV: VIEWING MULTI-LEVEL MAPS WITH THE MAP VIEW

Welcome to Mercury Application Mapping User's Guide

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

Using this Guide

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**

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

Part III   **Views, Organization Rules, and Notifications**

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 Multi-level Maps with the Map View

Describes how to use the Map View 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 object 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 View Indicators and Charts

Describes how to generate reports, view indicators, and charts

Part VI   Object Relationships

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

Part VII   Logical Objects and Rules

Describes how to use the Logical Object Builder Manager to create logical objects 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 Class Browser

Describes how to use the Class Browser to view the class model of your managed world, to create and add new classes to the class 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 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*.

Welcome

Part I

Introduction

1

Introduction to Mercury Application Mapping

This chapter provides a general introduction to Mercury Application Mapping.

This chapter describes:	On page:
Mercury Application Mapping System Architecture	5
Mercury Application Mapping Discovery Process	6
Discovery Process Architecture	6
Mercury Application Mapping Topology Query System	8
Mercury Application Mapping Deployment	13

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 objects 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

➤ **Root-cause analysis and problem resolution** – understanding the casual relations between objects 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 8). 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 "Defining Managed Views" on page 91. 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 objects 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:
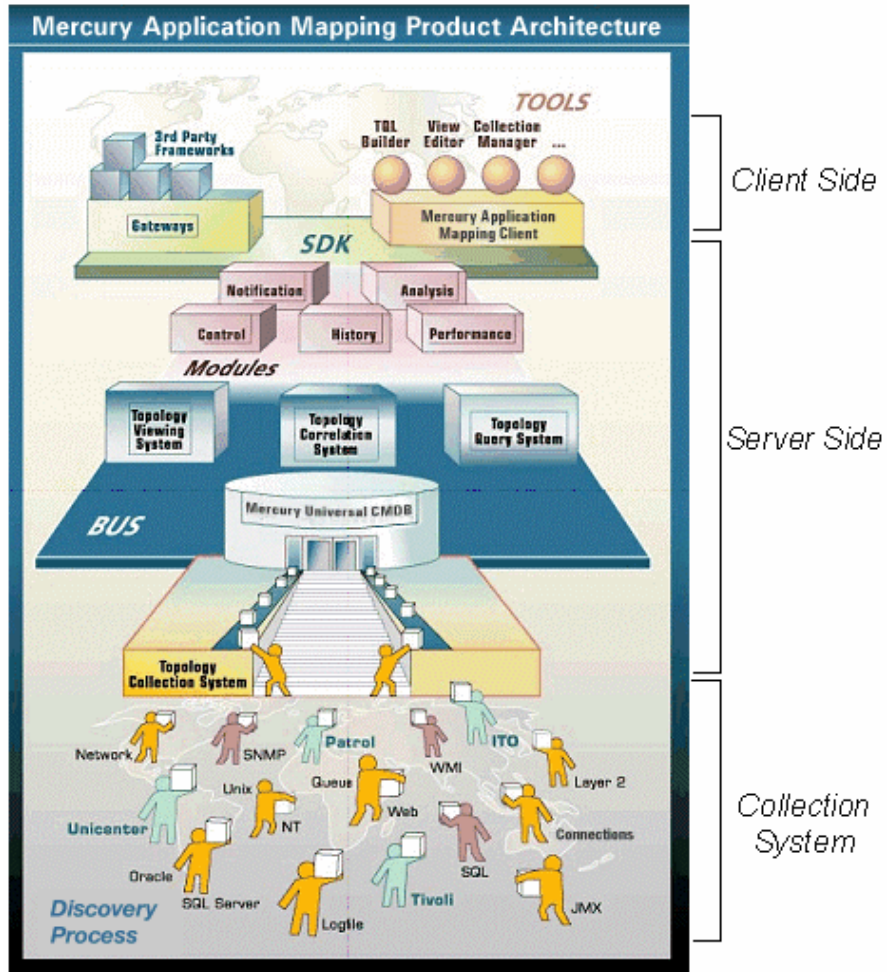
Mercury Application Mapping Discovery Process

The Mercury Application Mapping discovery process is the mechanism that discovers IT infrastructure resources and their interdependencies, and delivers the discovered data to a central repository. The IT resources that can be discovered through the discovery system include: applications, databases, network devices and different types of servers.

The Mercury Application Mapping discovery process is run by activating discovery patterns. A discovery pattern is an XML template, which 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 Mercury Universal CMDB, and the executing schedule. A discovery pattern contains a description of the objects and links that are created with each specific pattern. The object and link definitions are taken from the Class Model, which contains the definitions of all the classes and links defined in the system.

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. The objects that are first discovered are placed in the Mercury Universal CMDB. They act as triggers that activate a discovery pattern. Every time a discovery pattern is activated, it discovers more objects, which in turn are used as triggers for other discovery patterns. This process continues until your entire IT infrastructure is discovered and mapped.

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

Discovery Process Architecture

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

➤ "Mercury Application Mapping Server" on page 7

➤ "Probe Manager" on page 7

➤ "Probe Gateway" on page 7

Mercury Application Mapping Server

The Mercury Application Mapping server manages the discovery process. It delivers discovery requests, receives discovery results, and stores collected data in the Mercury Universal CMDB.

Probe Manager

The Probe Manager runs the discovery process. Its function is to receive a discovery task from the Probe Gateway, run the task according to the discovery pattern's schedule, and collect the data.

Probe Gateway

The Probe Gateway is a communication server that is located on the remote site. Only one Probe Gateway resides at each remote site, and it is the sole communication server for the entire domain. The Probe Gateway connects the Mercury Application Mapping server to the remote site, and provides a way of communication between the Mercury Application Mapping server and one or more Probe Managers.

When a discovery task is invoked, the Mercury Application Mapping server dispatches the task to the Probe Gateway. Each discovery task contains the name of the Probe Manager to where the discovery task should be sent. Once the Probe Gateway gets a task, it channels it to the appropriate Probe Manager. The Probe Gateway can store a task request when the appropriate Probe Manager is not available. Similarly, it stores task results if the connection with the Mercury Application Mapping server is lost. Once the Probe Manager completes a task, it passes the discovered data to the Probe Gateway. The Probe Gateway receives the data, and then forwards it to the Mercury Application Mapping server.

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 8).

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 Mercury Universal CMDB. TQL is used to create queries that retrieve business service data from the Mercury Universal CMDB, and display that data. TQL queries constantly search the Mercury Universal 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 9).

The Mercury Universal CMDB is the core information repository of Mercury Application Mapping. It contains the class 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 9).

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 10).

Topology Query Language (TQL)

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

➤ TQL enables Mercury Application Mapping to draw conceptual links between objects, which represent their actual interdependencies. Using pre-defined operators, the different types of interconnections that exist between objects 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 Mercury Universal CMDB.

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

➤ Constantly searches the Mercury Universal 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 class 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 objects and links is deposited, grouped and updated in the Mercury Universal CMDB in the form of class definitions according to object-oriented methodology.

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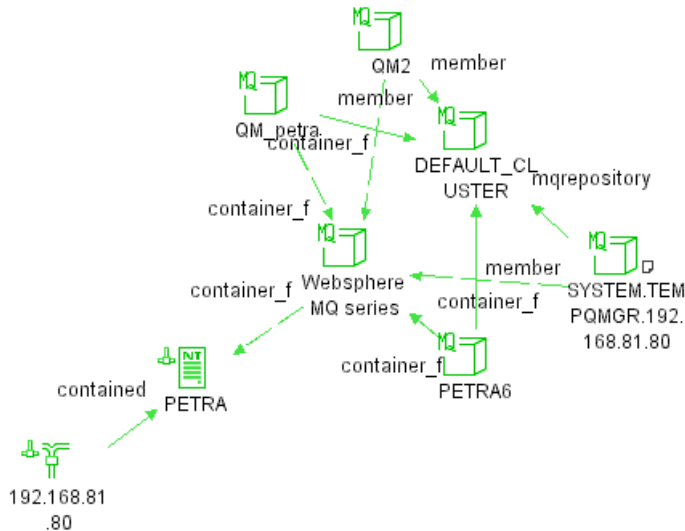
The nodes and links are organized in the Mercury Universal CMDB according to object-oriented methodology.

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 objects 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 Mercury Universal 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 Mercury Universal CMDB for changes in objects 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 Logical Objects:
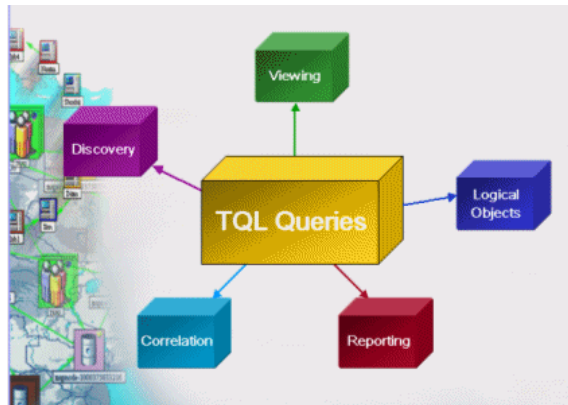
➤ **Viewing System (Mapping System)** – Generates specific managed views (maps) that display the query results and enable their monitoring and management. The Viewing System provides the user with three sets of rules for achieving the kind of display that serves his/her needs best:

➤ **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 objects 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 Rules to Links" on page 109.

➤ **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 an object in a Normal state, the object starts blinking and this blinking passes on to the layers above the object. The user can control this blink propagation and decide whether the blinking propagates and to which view layers. For details, see "About the Event System" on page 214.

➤ **Notification Rules** – For defining notifications to be sent automatically once changes occur in the availability, status, correlation, or attribute values of managed objects. For details, see "Sending Notifications" on page 113.

➤ **Correlation System** – Defines casual relationship between objects, which enables to locate and identify, in case of system's alerts, the object 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 object'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 335.

➤ **Logical Object System** – Enlarging the Mercury Universal CMDB, by adding objects and links 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 Logical Objects and Rules" on page 303.

Mercury Application Mapping Deployment

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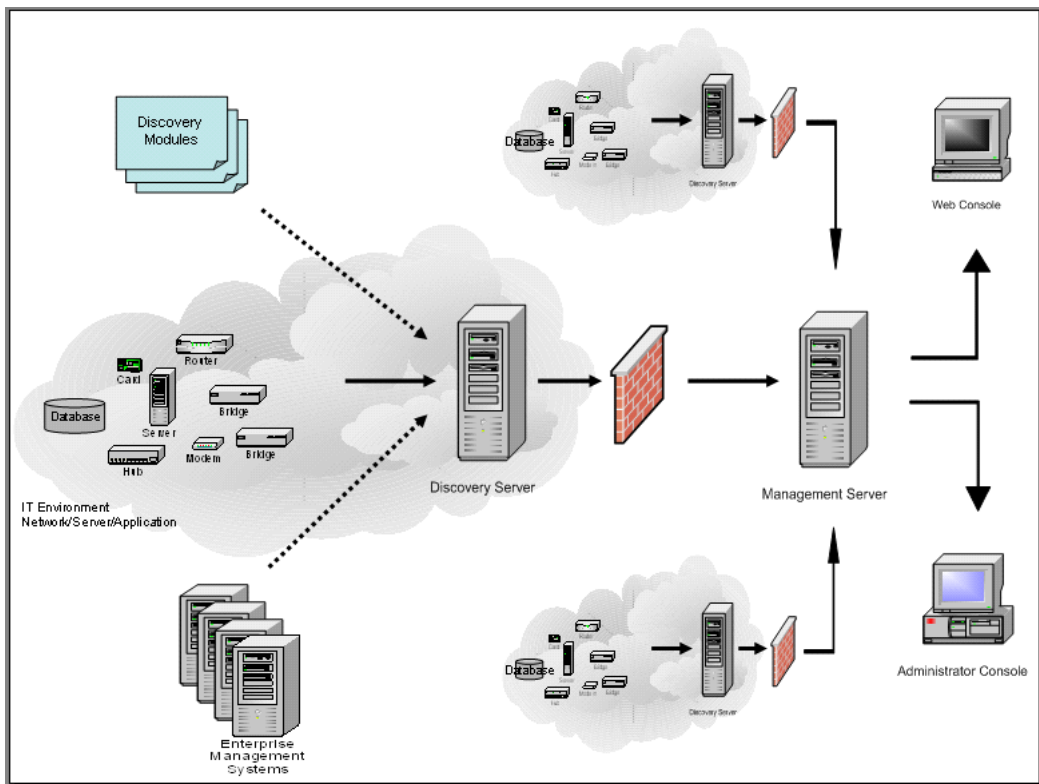
2

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
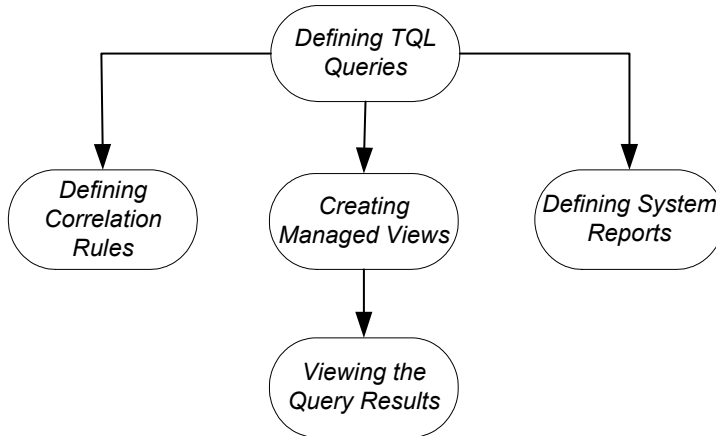
workflow summarizes the main tasks performed, as described below:

➤ **Define a new query** – You define a new query, using the Topology Query Language, which consists of selecting objects and defining the specific attributes and conditions of the objects and links that define the relationship between the objects. 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."

➤ **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 Map View. This includes adding organizational rules and notifications to define the organization structure of the objects in the Map View. For details, see Part III, "Views, Organization Rules, and Notifications."

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

➤ **Create and define map views and charts** – You create and define map views and charts, and produce view indicators that provide statistical information and visual representation of data. You can then navigate between different map levels to focus on the specific area containing your query results. For details, see Part V, "Generating View Indicators and Charts."

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

➤ **Create and define logical object rules** – Theses rules can be used for several purposes: creating new logical objects and links for adding data to the Mercury Universal CMDB and allowing additional representation options of existing data; updating the value of object attributes in the database, and deleting objects from the database. For details, see Part VII, "Logical Objects and Rules."

➤ **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."

➤ **Monitor changes in the managed data –** For details, see Part IX, "The Class Browser."

➤ **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

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.
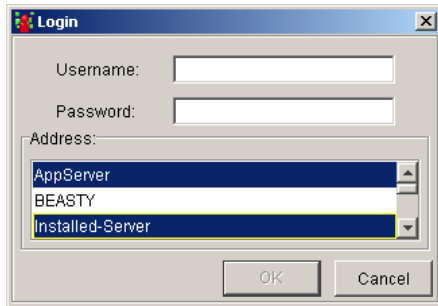
By displaying several servers in one window, you can view and manage them simultaneously and reuse TQLs you prepared for one server for other servers as well.

On the Map View, the different servers are represented by different icons. In other Mercury Application Mapping managers, the different servers are represented by different folders, as is shown in the following:

Two different servers that are managed in the View Manager
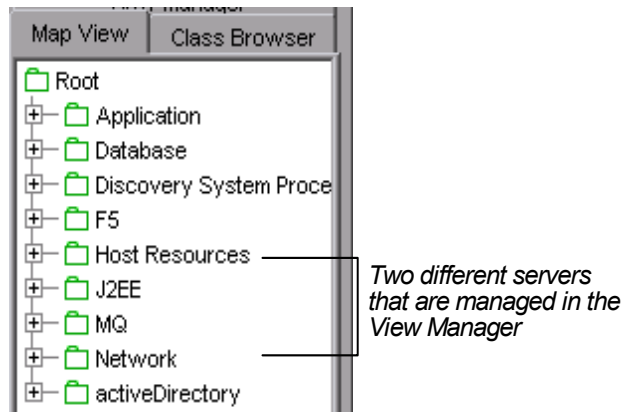
5 After you enter all the Login details, click **OK** to display the Mercury Application Mapping main window, as described in the next section.

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.
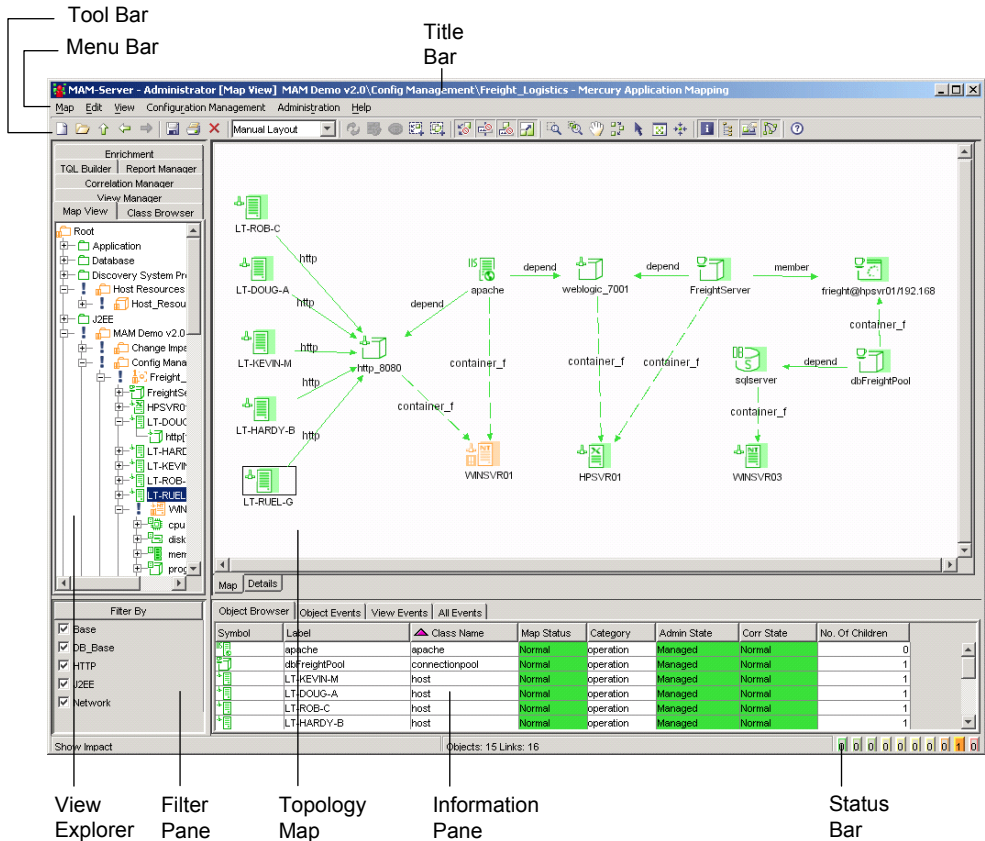
➤ **Title bar** – Displays the following information:

➤ 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.

➤ **Explorer pane** – Contains hierarchical tree structures of the class model, TQL queries, managed views or map objects, depending on the tab selected (for details, see "Explorer Pane" on page 22).

➤ **Filter pane** – Contains object groups that can be displayed or hidden in the Map pane.

➤ **Map pane** – Displays multi-layered views of objects and links, including the relationships between them, using various link types. Each class has a unique icon; its neighbors (as indicated by its connection to the class in the Explorer pane) all share the same shape surrounding the icon. (Shapes are predefined according to logical class groupings.) The Map pane displays different View Indicators depending on the View Indicators selected in the **View** menu. To hide or display the View Indicators, select **View** > **View Indicators** and click the View Indicator you want to hide or display.

In addition, the Map pane includes the following components:

➤ When the TQL Builder tab is selected, the class model pane is displayed.

➤ When the Map 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 23.

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

➤ **Status bar –** For details, see "Status Bar" on page 38.

This section contains the following topics:

Explorer Pane

The Explorer pane contains hierarchical tree structures of the class model, TQL Queries, managed views or map objects, 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 objects – nodes and physical links. 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.

➤ **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 View Manager" on page 83.

➤ **Map View** – Enables you to view the results of your TQL queries based on the view defined with the View Manager. The maps displayed in the Map View 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 objects. The Map View also enables you to view all events associated with the objects displayed in each map. For details, see "Introduction to the Map View" on page 133.

➤ **Class Browser** – Enables you to view the information in the class model, which consists of all the defined classes in the system and the links that define the relationships between them. The class model represents the structure of the managed world on which Mercury Application Mapping is used. For details, see "Introduction to the Class Browser" on page 365.

➤ **Enrichment** – 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 Logical Objects and Links" 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 329.

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 Map View):

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

➤ New views in the View Manager. For details, see "Creating a View" on page 92.

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

➤ New logical objects and links in the Enrichment manager. For details, see "Creating Logical Objects and Links" on page 306.

➤ New System Reports in the Report Manager. For details, see "Creating a Report" on page 337.

➤ New classes in the Class Browser. For details, see "Creating Classes" on page 374.

➤ **Open** – Enables you to open a selected parent object and drill down to its lower map layers in the Map View. For details, see Chapter 10, "Using Map View."

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

➤ TQL queries in the TQL Builder. For details, see "Saving a Query" on page 74.

➤ Managed views in the View Manager. For details, see "Saving the Managed View" on page 112.

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

➤ Any changes to the manual layout in the Map View. For details, see "Improving Views" on page 172.

➤ System Reports in the Report Manager. For details, see "Viewing and Saving the Generated Report" on page 345.

➤ **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 74.

➤ Managed views in the View Manager. For details, see "Saving the Managed View" on page 112.

➤ 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 329).

➤ Logical Object rules (for details, see "Introduction to Logical Objects and Links" 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 74.

  ➤ Managed views in the View Manager. For details, see "Saving the Managed View" on page 112.

  ➤ 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 329).

  ➤ Logical Object rules (for details, see "Introduction to Logical Objects and Links" on page 295).

  For details on exporting a set of TQLs, see "Exporting a Set of TQL Items" on page 75.

➤ **Layer Setup** – Enables you to accurately define the view's layer layout, by assigning positions for the nodes and links of the view, to achieve the clearest presentation possible. For details, see "Improving Views" on page 172.

➤ **Export Graph** – Enables you to save a Map View to file. For details, see "Saving a Map View" on page 196.

➤ **Change Password.** – Enables you to change your login password. For details, see "Changing Your Password" on page 43.

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

➤ **Print** – Enables you to print the contents of the Information pane in the Map View. For details, see "Printing the Map View" on page 198.

➤ **Print Map** – Enables you to print the contents of the Map pane in the Map View. For details, see "Printing the Map View" on page 198.

➤ **Exit** – Closes the application.

Edit Menu

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

➤ **Insert Object** – Displays the Insert Object dialog box, enabling the administrator to manually add objects to the database. For details, see the *Mercury Application Mapping Administration Guide*.

➤ **Add Background Image –** Enables you to add a background map for display and navigational purposes in the Map View. This is used with the Manual Layout option. For details, see "Adding a Background Image to a Map View" on page 174.

➤ **Filter Editor –** Enables you to create and use existing event filters to reduce the number of events displayed in the event tabs, browsers, and log. Only those events that meet the criteria defined by the selected filter are displayed, enabling you to focus on events of particular interest and significance. You can define multiple filters, but only one filter can be active at a time. For details, see "Setting and Using Event Filters" on page 229.

➤ **Find –** Enables you to search for objects in the Map View, 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 objects according to their class, label and/or attributes. For details, see "Searching for Objects" on page 202

View Menu

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

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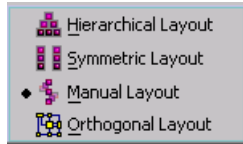
➤ **Hierarchical Layout** – Enables you to display the precedence relations in the topology:
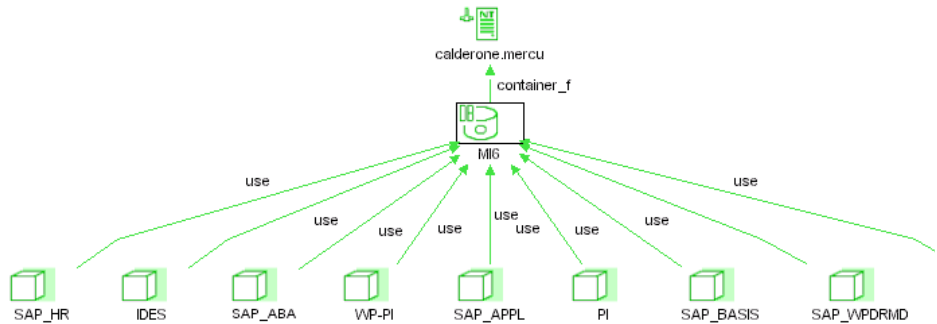
➤ **Symmetric Layout –** Enables you to display clear representations of complex networks:
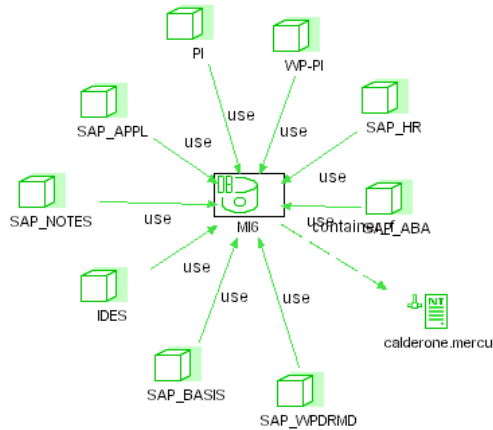
➤ **Manual Layout** – Enables you to manually change the Map View display, by dragging the displayed nodes and links to your required location.
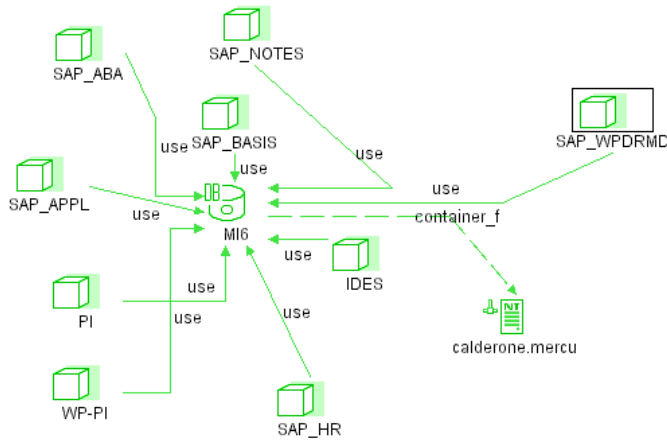
➤ **Orthogonal Layout** – Enables you to display views of outstanding clarity that are achieved by employing only horizontal and vertical edge routing.
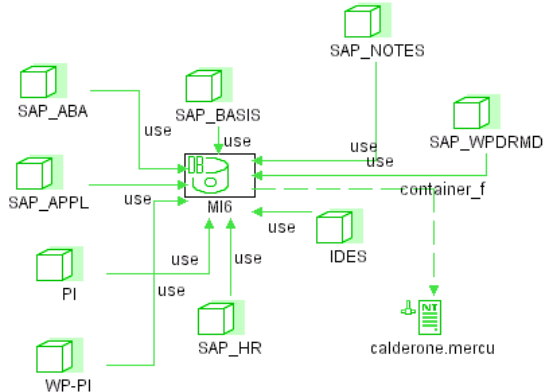
Note: You can save the manual changes you made, by clicking the **Save** button on the toolbar. Thus, you are able to reuse the layout you created.

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 map view 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.

➤ **View Indicators –** Enables you to generate and display dynamically changing diagrams, which contain statistical data about several significant attributes of the system's views. Each View Indicator relates to a specific attribute(s) of the view or the view's objects. When it is generated, it appears in a new and separate tab on the Map pane, and displays the results in a fixed diagram type (pie, bars or dashboard). For details, see Chapter 14, "Generating View Indicators".

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

  ➤ "Object Browser Tab" on page 137

  ➤ "Object Events" on page 219

  ➤ "View Events" on page 219

  ➤ "All Events" on page 219

  ➤ "Inbox" on page 219

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

➤ **Details –** Enables you to display the attributes of all objects that are located in the same view layer in one table. Data is organized by object class. For details, see "Viewing Layer Details" on page 178.

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 Map pane.

Configuration Management

The **Configuration Manager** 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 "Comparing Compound Objects" on page 421

➤ **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 345.

➤ **Gold Master Reports –** Enables you to generate a report that compares the configuration of a Gold Master object to other objects of the same class (for details, see "Gold Master Comparison" on page 411.

➤ **Impact Analysis** – **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 "Impact Analysis" on page 441..

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 the *Mercury Application Mapping Administration Guide*.

➤ **User Management** – Displays the following:

  ➤ **User Profile –** Displays the User Profile dialog box, enabling you to display in the Map View 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 in *Mercury Application Mapping Administration Guide*.

  ➤ **User Manager** – Displays the User Manager dialog box, enabling you to create and manage the system's users. For details, see the *Mercury Application Mapping Administration Guide*.

  ➤ **Role Manager** – Displays the Roles dialog box, enabling you to create and manage roles. For details, see 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, see 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 Map View or in other Managers, except for the Class Browser, 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 Class Browser, 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 an object's management state. For details, see 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 Chapter 26, "Enumerations and Lists."

➤ **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 the *Mercury Application Mapping Administration Guide*.

Help Menu
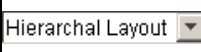
The **Help** menu contains the following options:
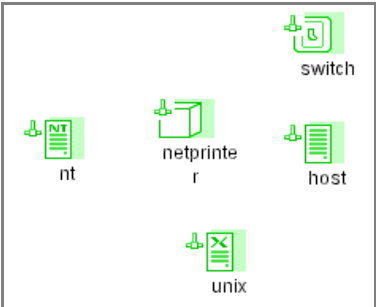
➤ **Mercury Application Mapping Help** – Opens the Documentation Index.

➤ **About Mercury Application Mapping** – Displays copyright and 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	Icon	Use This Option to...
New		Define the following: ➤ New instance view in Map View. For details, see "Creating a New Instance View" on page 207. ➤ New queries in the TQL Builder. For details, see "Creating a TQL Query" on page 59. ➤ New views in the View Manager. For details, see "Creating a View" on page 92. ➤ New correlations rules in the Correlation Manager. For details, see "Creating a Correlation Rule" on page 284 ➤ New logical objects and links in the Enrichment Manager. For details, see "Creating Logical Objects and Links" on page 306. ➤ New System Reports in the Report Manager. For details, see "Creating a Report" on page 337. ➤ New classes in the Class Browser. For details, see "Creating Classes" on page 374.
Open		Open a selected parent object and drill down to its lower map layers in the Map View. For details, see "Map View shows the results of the currently selected view in the View Explorer, and consists of objects and links that are defined in the view. You can change the way the maps are displayed by selecting one of the view layouts (Hierarchical, Symmetric, Manual, or Orthogonal). For details, see "Understanding Layer Layouts" on page 154. You can also modify the values of the default layer layouts, as described in "Defining a Layer Layout" on page 173." on page 152.
Up		Move up one level in the Map View. For details, see Part IV, "Viewing Multi-level Maps with the Map View."
Back		Return to the last view you displayed.

Option Name	Icon	Use This Option to...
Forward		Display a view you displayed before clicking the **Back** button, click the **Forward** button.
Save		Save new and changed items in all Mercury Application Mapping Managers (apart from the Class Browser):
		➤ TQL queries in the TQL Builder. For details, see "Saving a Query" on page 74.
		➤ Managed views in the View Manager. For details, see "Saving the Managed View" on page 112.
		➤ Correlations rules in the Correlation Manager. For details, see "Saving the Correlation Rule" on page 291.
		➤ Any changes to the manual layout in the Map View. For details, see "Improving Views" on page 172.
		➤ System Reports in the Report Manager. For details, see "Saving the System Report Template" on page 344.
Print		Print the contents of the Map pane and the contents of the Information pane in the Map View. For details, see "Printing the Map View" on page 198.
Delete		Delete a selected item (such as a view or TQL) from the TQL Builder or View Manager.
Layout	Hierarchal Layout	Select the layout display from the list. For further details about each of the options, see "View Menu" on page 27.
Refresh		Reorganize the objects according to the selected layer. This is used when objects have been moved around or when hidden objects exist (used mainly for the Map View).
Rebuild		Rebuild the view in the Map View (used when the view name is selected in the tree).
Show Hidden Symbols		Display previously hidden objects.
Select All		Select all the objects and links in a selected layer.

Option Name	Icon	Use This Option to...
Select All Nodes		Select all the nodes in a selected layer.
Hide All Links		Toggle between showing or hiding the links in the displayed view.
Hide All Link Labels		Toggle between showing or hiding the link 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 switch nt netprinter host unix
Zoom		Zoom in on a specific section of the map view 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 map pane.
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.

Option Name	Icon	Use This Option to...
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 links.
Select Mode		Select an object or objects.
Fit to Window		Fit the complete view in the window.
Centers the Selected Object in Graph		Places the selected object in the center of the map.
Information Pane		Display or hide the Information pane (containing the Object Browser, Object Events, View Events, and All Event tabs in the Map View).
Explorer Pane		Display or hide the Explorer pane in the selected tab.
Filters Pane		Display or hide the Tags pane in the selected tab.
Map Pane		Display or hide the Map pane in the selected tab.
Help Topic		Display the table of contents for the online Help.

Status Bar

The status bar at the bottom of the Mercury Application Mapping main window displays the number of nodes and links displayed in the selected (or current) layer, as well as the number of current active events per severity level.

The Main Window's Level Structure

Mercury Application Mapping Main Window has four basic levels: Server, Folder, View, and Object.

The Folder, View, and Object levels can include additional sublevels or layers. The Server and Folder levels may not appear in the main window – their appearance depends on the number of servers to which you are connected and on the internal organization of views in the folders.

To drill down from a higher level to a lower one, double-click the symbol on the Map pane. For more navigation options, see "Map View shows the results of the currently selected view in the View Explorer, and consists of objects and links that are defined in the view. You can change the way the maps are displayed by selecting one of the view layouts (Hierarchical, Symmetric, Manual, or Orthogonal). For details, see "Understanding Layer Layouts" on page 154. You can also modify the values of the default layer layouts, as described in "Defining a Layer Layout" on page 173." on page 152.

This section includes:

➤ "Server Level" on page 39

➤ "Folder Level" on page 40

➤ "View Level" on page 41

➤ "Object Level" on page 42

Server Level

If you chose to connect to several servers during log in, the top level of the Map view shows the servers to which you are connected:
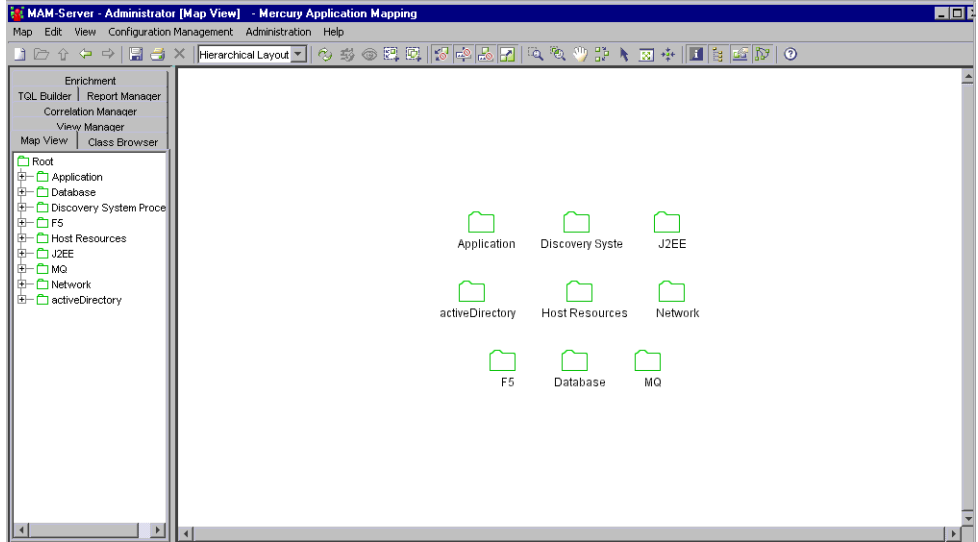
Folder Level

If you choose to connect to one server only 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:
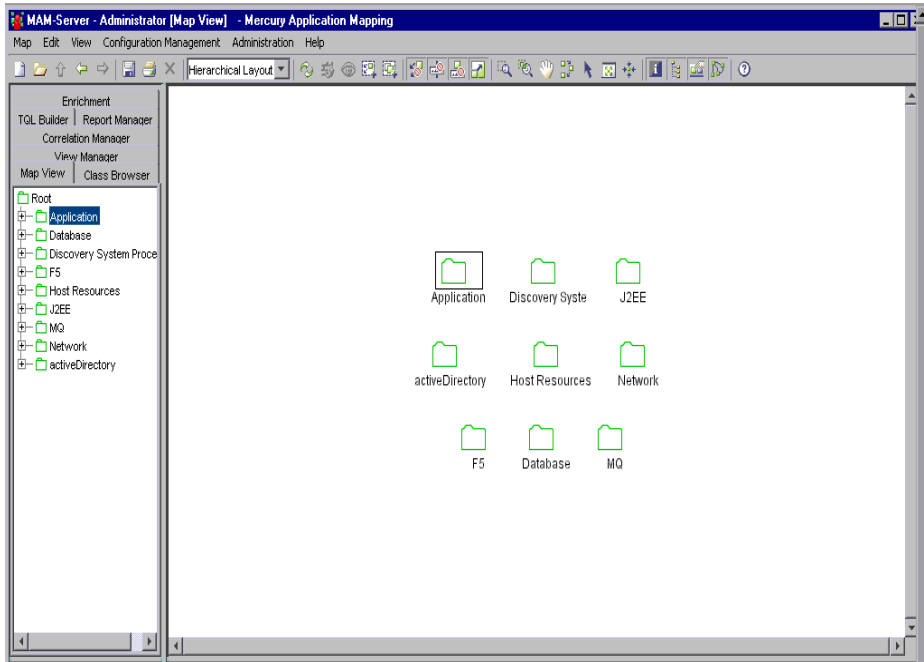
A folder can contain subfolders and several views. You use the View Manager to organize the views in a folder. For details, see "View Shortcut Menu" on page 86.

View Level

The View level displays the views that are set in the View Manager to be displayed on the Map View, and the views that are part of your user profile. For details, see "Understanding Layer Layouts" on page 154.
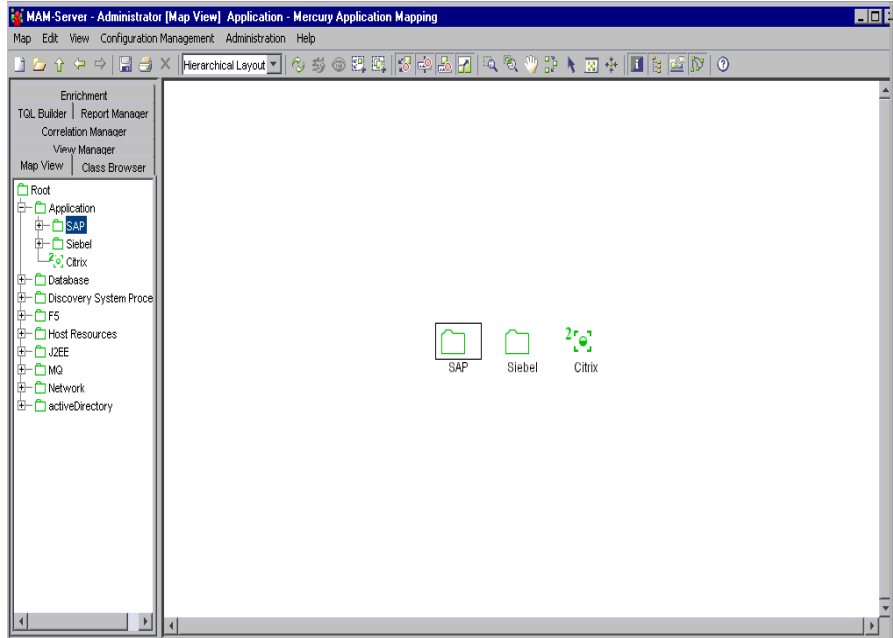
Note: A view's icon can be customized in the View Manager. For details, see "Map View shows the results of the currently selected view in the View Explorer, and consists of objects and links that are defined in the view. You can change the way the maps are displayed by selecting one of the view layouts (Hierarchical, Symmetric, Manual, or Orthogonal). For details, see "Understanding Layer Layouts" on page 154. You can also modify the values of the default layer layouts, as described in "Defining a Layer Layout" on page 173." on page 152.

Object Level

The Object level displays the view results, that is, the objects and links that fulfill a View TQL query. These objects and links are displayed according to display and organizational rules that are assigned to them in the View Manager:
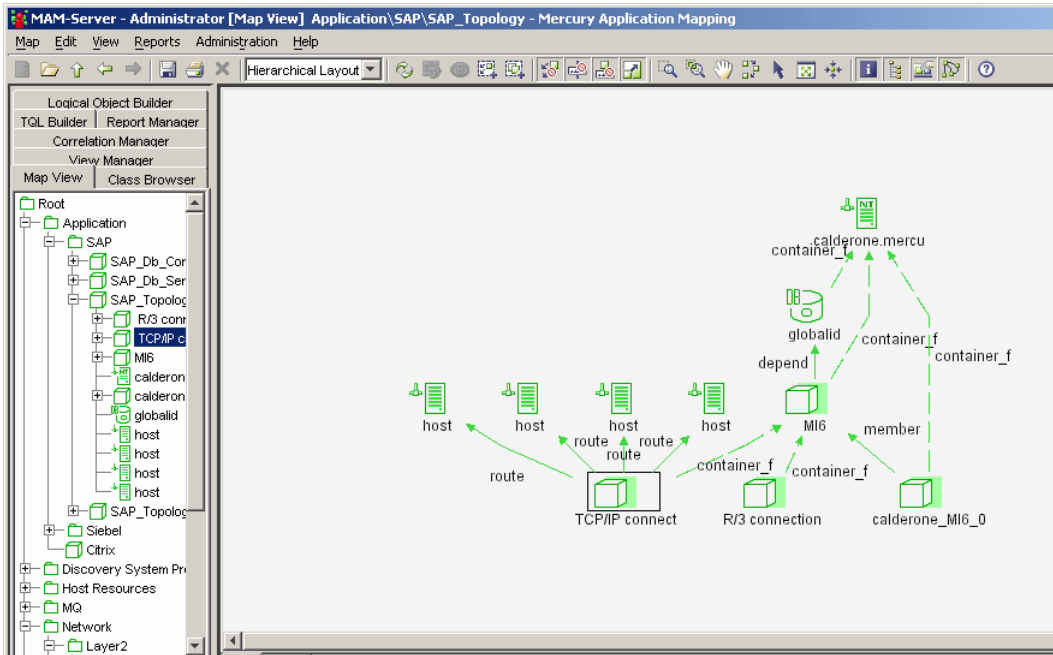
The icon that represents each object and link is determined by the object/link class. Object icons are made up of the following three components:

 ➤ 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 object 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:

➤ "Changing Your Password" on page 43

➤ "Selecting Multiple Objects" on page 44

➤ "Working with Shortcut Menus" on page 44

➤ "Tooltips" on page 44

➤ "Drilling Down" on page 44

➤ "Operator List" on page 45

Changing Your Password

This section describes how to change the password with which you log in to Mercury Application Mapping.

To change your password:

1 From the **Map** menu, select **Change Password** to open the Change Password dialog box.

2 In the **Old Password** box, type the old password.

3 In the **New Password** box, type the new password.

4 In the **Confirm New Password** box, type the new password again.

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

Selecting Multiple Objects

In the Map pane you can select more than one object.

To select multiple objects:

➤ Hold down **Ctrl** and click the objects or drag the mouse diagonally along the objects as shown in the following figure:

Before   After
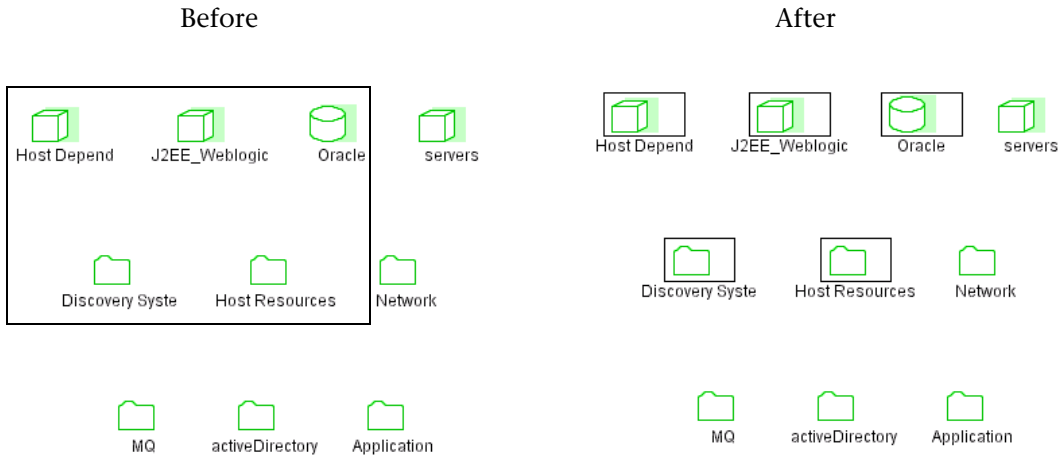
Working with Shortcut Menus

Right-clicking an object in the Map pane opens a shortcut menu. The shortcut menu depends on the selected object. The shortcut menus are detailed in each relevant part of this guide.

Tooltips

When the cursor is moved over an object in the Map pane, a tooltip displays information that is relevant to the object.

Drilling Down

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

Operator List

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

Operator	Description
=	Checks whether the attribute value is equal to the value specified in the value box.
= (ignore case)	Checks whether the attribute value is equal to the value specified in the value box independently of the case.
!=	Checks whether the attribute value is not equal to the value specified in the value box.
>	Checks whether the attribute value is greater than the value specified in the value box.
>=	Checks whether the attribute value is greater than or equal to the value specified in the value box.
<	Checks whether the attribute value is less than the value specified in the value box.
<=	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.
NOT LIKE	Uses a wildcard (%). The same as 'Like' but looks for attribute values that do not include the string.
IN	Displays only the instances where this attribute value equals one of the selected values. For example, for objects 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.
NOT IN	Displays only the instances where this attribute does not equal one of the selected values.
IS NULL	Checks whether the attribute value is null.
IS NOT NULL	Checks whether the attribute value is not null.

Operator	Description
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

3

Introducing the Topology Query Language (TQL) Builder

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

This chapter describes:	On page:
Working with TQL Builder	51
TQL Builder Shortcut Menus	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 persists in the system memory and generates updated results automatically.
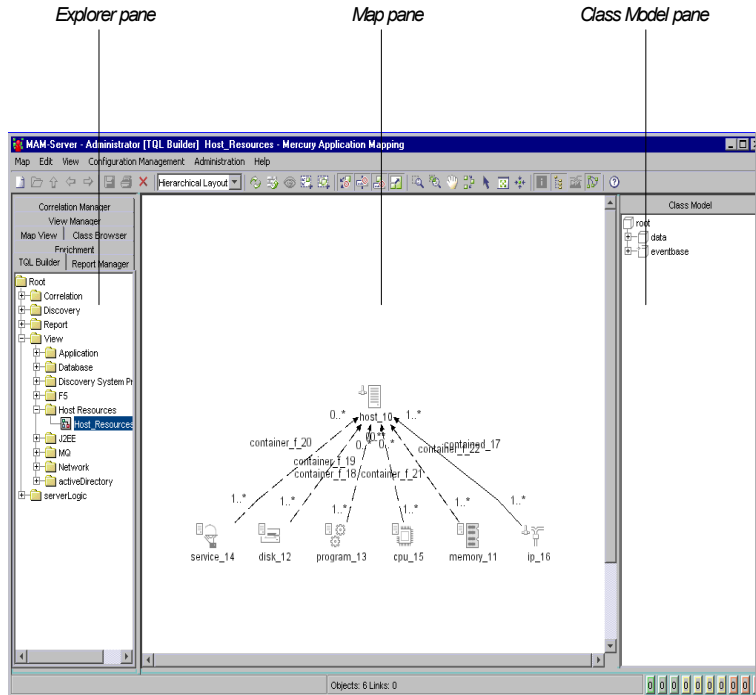
When the TQL Builder tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **Explorer pane** – 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.

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

➤ **Class Model pane** – Represents the class model and contains icons for each object, as defined by the administrator (for details, see "Assigning an Icon to a Class" on page 378). By clicking and dragging objects to the Map 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** objects to the Map pane and then define the relationship between them by adding links.

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 objects and links in the Map pane.

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

➤ Define the specific attribute conditions of objects and links in the query.

➤ Create and delete tags, which are logical groupings of objects and links in the query.

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

➤ Define whether to display the result of the query in the Map View, including the objects derived from the selected object.

➤ 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 Map pane to view the selected TQL at different zoom levels.

➤ Display the queries in different layout views.

TQL Builder Shortcut Menus

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

➤ "TQL Folder Shortcut Menu" on page 52"

➤ "TQL Query Menu" on page 53

➤ "Multiple Object Menu" on page 54

➤ "Single Object Shortcut Menu" on page 54

➤ "Link Shortcut Menu" on page 55

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 Explorer pane:

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.
Rename	Enables you to edit the label of a selected 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 Explorer pane:

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 74.
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 61.
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: Right-clicking an empty area in the Map pane when a TQL is displayed, opens a shortcut menu with zoom and layout options, similar to those available from the **View** menu. For details, see "Working with Shortcut Menus" on page 44.

Hold the cursor over a node or link to view a tooltip. The tooltips for the nodes and links contain the definitions of the selected objects, for example, the attribute conditions, as described in "Defining Node and Link Attribute Conditions" on page 70.

Multiple Object Menu

The following table provides a brief description of each option in the shortcut menu displayed by selecting two or more objects in the Map pane. For a description of how to select multiple objects, see "Selecting Multiple Objects" on page 44.

Option	Description
Delete	Enables you to delete the selected objects.
Add Link	(For two selected objects only) Displays the Add Link dialog box, enabling you to add a predefined link between two objects in a query, as described in "Adding Nodes and Links to a Query" on page 66. (Only relevant links are displayed on the **Add Link** list.)
Add General Links	Displays the General Link dialog box, enabling you to add multiple links at the same time to a query, as described on "Adding Multiple Links" on page 70.
Copy	Enables you to copy the selected objects.
Paste	Enables you to paste the copied objects.

Single Object Shortcut Menu

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

Option	Description
Delete	Enables you to delete the selected object.
Add Link	Displays the Add Link dialog box, enabling you to create an internal link for your object by selecting it from a predefined list, as described on "Adding Nodes and Links to a Query" on page 66.
TQL Node Definition	Displays the TQL Node Definition dialog box, enabling you to define the attribute conditions for the selected object, as described on "Defining Node and Link Attribute Conditions" on page 70.

Option	Description
Condition Junction Definition	Allows you to define link conditions under which the link's connecting node will be included in the TQL query results. For details, see
Copy	Enables you to copy the selected objects.
Paste	Enables you to paste the copied objects.

Link Shortcut Menu

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

Option	Description
Straight	(Enabled only for links with angles on the Map pane.) Enables you to straighten links with angles.
Delete	Enables you to delete the selected link.
TQL Link Definition	Displays the TQL Link Definition dialog box, enabling you to define the attribute conditions for the selected link, as described in "Defining Node and Link Attribute Conditions" on page 70.
Copy	Enables you to copy the selected objects.
Paste	Enables you to paste the copied objects.

4

Defining Topology Query Language (TQL) Queries

This chapter explains how to define TQL queries.

This chapter describes:	On page:
TQL Query Workflow	58
Creating a TQL Query	59
Adding Nodes and Links to a Query	61
Defining Link Conditions	68
Adding Multiple Links	70
Defining Node and Link Attribute Conditions	70
Saving a Query	74
Deleting an Attribute Condition	74
Importing a TQL Query	74
Exporting a Set of TQL Items	75

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 objects that are to be part of the query. You then define specific attribute conditions for each object, including the attributes of the links that define the relationships between objects.

Note: A TQL query is subject to certain validation restrictions (for details, see "Validation Restrictions" on page 77)

TQL Query Workflow

You create TQL queries according to the following workflow:

➤ Create a new TQL query.

For details, see the next section.

➤ Add nodes and links to the query.

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

➤ Define node and link attribute conditions.

For details, see "Defining Node and Link Attribute Conditions" on page 70.

➤ Save the new query.

For details, see "Saving a Query" on page 74.

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 61

Creating a New Query

You can create a new query.

To create a new query:

1 Click the **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 Map View. For details, see Part IV, "Viewing Multi-level Maps with the Map View."

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

➤ **Logical Object** – Create queries regarding logical objects and links, whose existence can be deduced from prior or accumulated knowledge about your system. For details, see Chapter 19, "Defining Logical Objects and Rules."

> ➤ **Report** – Creates queries whose results are compiled into reports using the Report Manager. For details, see Chapter 21, "Introduction to the Report Manager."

> ➤ **Discovery** – Creates queries for use by the discovery system.

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, Logical Object and Report TQLs to be valid, they must comply with the restrictions described in "Validation Restrictions" on page 77. 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 Explorer pane.

Note: You can change these options at any time by right-clicking the query and selecting **TQL Properties** from the TQL shortcut menu (for details, see "TQL Builder Shortcut Menus" on page 52).

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 query:

1 On the Explorer pane, 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.

Adding Nodes and Links to a Query

After you have created the TQL query, the next step is to add the nodes and links that define the query. The nodes represent the objects, as defined in the class model, and the links represent the relationships between them. Links are defined one at a time for each pair of nodes in the query. Alternatively, you can use general links (for details, see "Adding Multiple Links" on page 70) to select a predefined series of relationships among multiple objects. For details, see "Adding Multiple Links" on page 70.

This section contains the following topics:

➤ "Functional Links" on page 62

➤ "Physical/Logical Links" on page 64

➤ "Adding Nodes and Links to a Query" on page 66

Functional Links

Functional links describe a physical or logical connection between two classes, which is not represented in the Mercury Universal CMDB. Each of the functional links uses a different function to create the connection between the classes.

The functional links are:

➤ **join_f** – Connects two objects that have at least one comparable value. When creating this link, you must define an attribute for each object, whose value is used for comparison.

For example, you can compare the values of IP net address and Network net address, and Host DNS and IP address. For details on defining **join_f** conditions, see step 6 in "Adding Nodes and Links to a Query" on page 66.

➤ **host_f** – Connects a host to an object that is fully dependent on it. The object can be any one of the host's lower-level descendents (children, grandchildren, and so forth).

For example, when the following connections exist: Host > SQL Server > SQL Database, you can connect the host directly to the SQL Database (the host's grandchild) using the **host_f** link (as shown in the following figure):
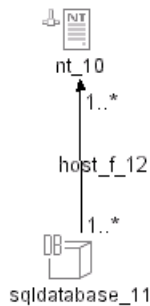
➤ **container_f** – Connects a host to an object that is fully dependent on it. The object must be one of the host's children.

For example, when the following connections exist: Host > SQL Server > SQL Database, you can add **container_f** links (parent-child links) between the host and the SQL Server and between the SQL Server and the SQL Database to establish an indirect connection between the host and the object (as shown in the following figure).

Tip: If you use **container_f** links, Mercury Application Mapping calculates the query results faster than if you use **host_f** links.

Physical/Logical Links

The following table contains the list of physical/logical links and their descriptions:

Links	Description	Usage	
Management	Represents a situation in which one object is managed by another.	Can connect between the following objects:	
		discoverymanager	discoveryprobegateway
		discoveryprobegateway	discoveryprobemanager
		discoveryprobemanager	discoverytask
parent	Describes the hierarchy between two objects in the Object Model.	Can connect between the following objects:	
		IP	interface
		IP	interfaceindex
		interfaceindex	interface
		port	interfaceindex
		exchangeserver	exchangesite
		exchangeroutinggroup	exchangesite
		exchangeserver	exchangeroutinggroup
		exchangeconnector	exchangeroutinggroup
		logicaldisk	disk
		printq	printer
		printq	printq
		logfile	logdir

Links	Description	Usage	
depend	Represents the dependency of one object on another.	Can connect between the following objects:	
		database	service
		service	sqlserver
		sqlfile	disk
		dbsnapshot	dbjob
		discoverypattern	discoverytql
		weblogic	ip
		weblogic	ipport
		connectionpool	database
		database	daemon
		service	service
member	Represents a situation in which two objects are members of the same group.	Can connect between the following objects:	
		exchangeserver	exchangeroutinggroup
		citrixserver	citrixfarm
		weblogic	j2eedomain
		ip	network
		host	network
contained	Used only to connect an ip to a host.	Can connect between the following objects:	
		ip	host
clientserver	Used only to connect an ipclient to an ipserver.	Can connect between the following objects:	
		ip client	ipserver

Links	Description	Usage	
tcp	Used for a TCP network connection when Mercury Application Mapping cannot determine which side of the IP connection is the ipclient and which side is the ipserver.	Can connect between the following objects:	
		ipunknown	ipunknown

Adding Nodes and Links to a Query

You can add nodes and links to a query.

To add nodes and links to a query:

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

2 From the tree displayed in the class model, click and drag one or more objects to the Map pane. These are the nodes that are included in the query.

3 Select two nodes by holding down **Ctrl** and clicking the nodes.

4 Right-click and select **Add Link** to open the Add Link dialog box.

The Add Link dialog box is displayed with the valid links that can be used between the selected nodes.

5 Select the link type that defines the relationship between the two nodes from the list.

6 If you did not select **join_f** continue to step 8.

7 If you selected **join_f**:

➤ Click **Next** to open the Conditions page.

➤ Click **Add** to open the Conditions dialog box:

- From the list of attributes for each node that you selected when you created the join_f link, select the attributes whose values you want to compare.

- From the **Operator** list, select the required operator (for details, see "Operator List" on page 45).

- Click **OK**.

➤ To define several conditions for one **join_f** link, repeat this step.

8 Click **Finish**. The nodes in the Map pane are connected by the new links you have defined.

The direction of the link 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.
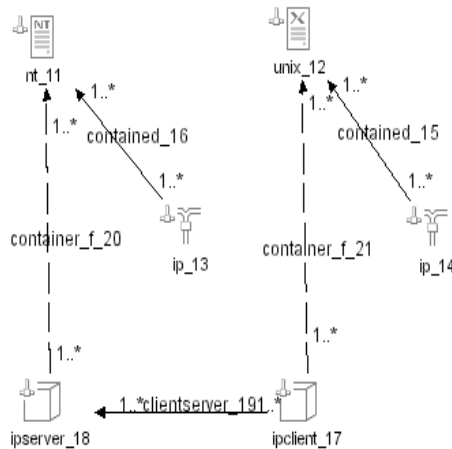
Tip: Pointing to a node or link displays a tooltip with the node's or link's tags and attribute conditions.

Copying/Pasting Objects in the TQL Builder

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

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

1 Select the object you would like to copy. You can select multiple objects as well.

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

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

Defining Link Conditions

Mercury Application Mapping enables you to define link conditions that allow the link's connecting node to be included in the TQL query results. You use the Condition Junction Definition dialog box to create an expression defining that condition.

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

The following examples of link conditions are all based on the sample TQL query shown below.
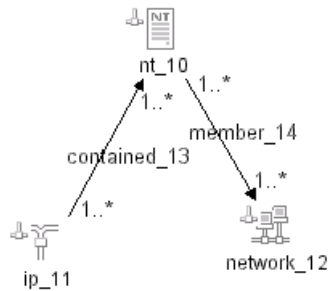
➤ [contained_13] OR [member_14] – means that the host must either have between two and seven IPs OR be a member of the network.

➤ [contained_13] AND [member_14] – means that the host must have between two and seven IPs AND ALSO be a member of the network.

➤ [contained_13] AND NOT [member_14] – means that the host must have between two and seven IPs AND NOT be a member of the network.

➤ [contained_13] OR NOT[member_14] – means that the host must have between two and seven IPs OR NOT be a member of the network.

To define a link condition:

1 In the TQL Builder, right-click any of the nodes in the selected TQL query and click to open the Condition Junction Definition dialog box. The dialog box contains a list of the links defined in the TQL.

2 Select the link or links that will be part of the definition by clicking the plus button to the left of each link. The selected link(s) is displayed in the **Expression** box (for details, see "Formatting Text" on page 101).

3 Click **Finish** to save the expression you have defined.

Adding Multiple Links

The following procedure for adding multiple links is optional.

To add multiple links to two or more selected nodes:

1 Select the required nodes to which you want to add multiple links.

2 Right-click any one of the selected nodes and select **Add General Links** to open the General Links dialog box.

3 Select the check box next to each link that you want to add to the query and click **OK**. Mercury Application Mapping automatically assigns the links to their proper places among the nodes and displays the results in the Map pane.

Note: General links cannot be used when selecting a single object.

Defining Node and Link Attribute Conditions

After you have added the nodes and links required for your query, you can define their specific attribute conditions. In addition, you can define whether to display the selected object in the Map View, and whether to display the objects derived from the selected object (that is, its neighbors) in the Map View.

This section includes:

➤ "Defining Attribute Conditions for Nodes" on page 71

➤ "Defining Link Conditions" on page 72

Defining Attribute Conditions for Nodes

You can define attribute conditions for nodes.

To define attribute conditions for nodes:

1 In the Map pane, right-click the object (node or link) whose conditions you wish to define and select **TQL Node Definition** or **TQL Link Definition**. The TQL Node/Link Definition dialog box is displayed.

2 Click **Add** to open the Condition dialog box.

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

4 From the **Operator** list, select the required operator (for details, see "Operator List" on page 45)

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

Attribute conditions cannot be defined for functional links. When using the operator **CHANGED_DURING**, you can only enter integers in the **Value** box. Integers, by default, are measured in minutes. You can change this default in the **appilogConfig.properties** file.

5 In the **Value** box, enter or select the value of the attribute. (The box changes according to the attribute type you selected.)

6 Click **OK**.

7 In the Attribute Condition tab, select **And** or **Or** to link multiple conditions.

8 Select the **Visible** tab:

➤ To display both the selected object and any objects derived from the selected object (inheritance objects) in the Map View, select **Derived**.

➤ To display the selected object in the Map View, select **Visible**. When
  Visible is cleared, an invisible box appears to the right of the selected
  object in the Map pane:
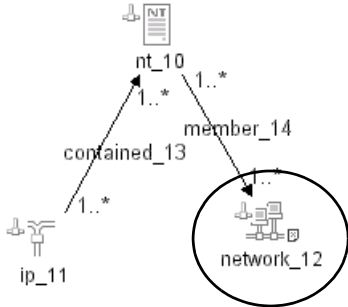
Any query results pertaining to that object are not displayed in the Map
View. This can be useful when certain links or nodes are required to build
the query but are not needed in the results. For example, NTs are
connected to the network by defining specific IPs, but in any query
results you may want to view the IP elements only and not the NT
elements.

Defining Link Conditions

You can define link conditions.

For example, if <end_1> is **ip** and <end_2> is **nt**, entering **1** in the first **Min**
box and asterisk (*****) in the first **Max** box instructs the system to include only
those **ips** that are connected to at least one **nt**. (The asterisk indicates an
infinite value.) Entering **3** in the second **Min** box and asterisk (*****) in the
second **Max** box instructs the system to include only those **nts** that are
connected to at least three **ips**.

To define link conditions:

1 Right-click the link and select **TQL Link Definition**.

2 Select the **Link** tab to display two sets of **Min** and **Max** boxes, one for each connection formed by the link (<end_1> > <end_2>, <end_2> > <end_1>).
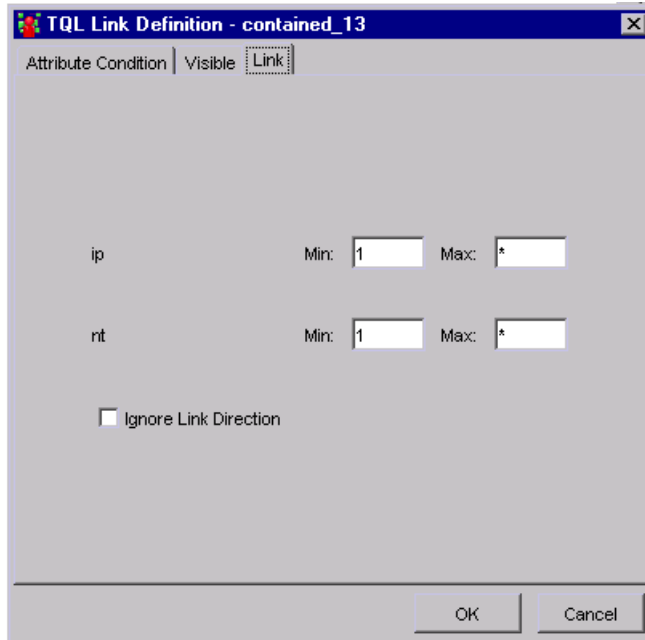
3 In both sets of **Min** and **Max** boxes, enter values that define the lower and upper limits for including the selected nodes in the query results.

4 Click **OK** to save the link or node definitions to the database. The new query is displayed in the Explorer pane.

Saving a Query

After you have defined the specific attribute conditions of the nodes and links in the query, the last step is to save the query to the database. 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 Map pane, and click **Save** on the toolbar or select a different query or tab in the Explorer pane. A message is displayed asking if you want to save the current query.

The query you created is saved to the Mercury Universal CMDB.

Deleting an Attribute Condition

You can delete an attribute condition

To delete an attribute condition:

1 In the Attribute Condition tab (for details, see "Defining Node and Link Attribute Conditions" on page 70), select the attribute condition you want to delete.

2 Click the **Delete** button or press the **Delete** key.

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 Explorer pane.

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:
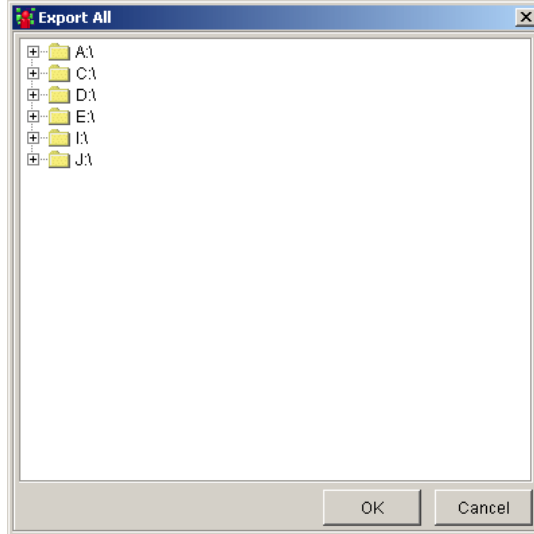
2 Browse and select the folder to which you want to export the TQL items
displayed in the Manager.

3 Click **OK**.

5

Validation Restrictions

This chapter explains TQL validation restrictions.

This chapter describes:	On page:
Correlation TQL Validation	79
Logical Object and Report TQL Validations	79

About Validation Restrictions

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

➤ **Internal Links** – A TQL must not contain internal links, that is, a link must not lead from a node to itself, as the following example illustrates:
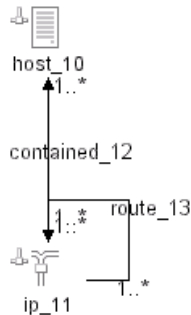
➤ **Circular Structure –** The TQL structure cannot be a closed circle, as shown in the following example:
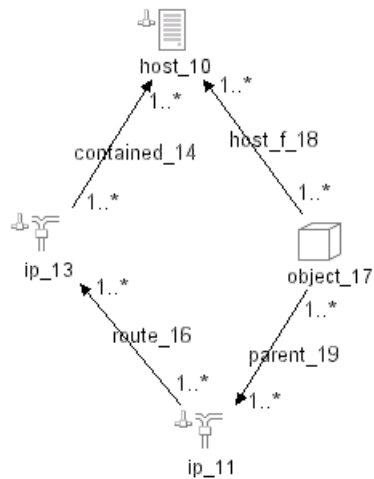
Note: There is one exception to the above rule: the TQL can be circular if one of its links is invisible. To define a node/link as invisible, see step 8 in "Defining Node and Link Attribute Conditions" on page 70.

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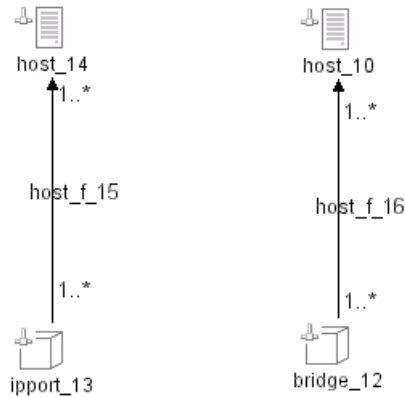
Correlation TQL Validation

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.

Logical Object and Report TQL Validations

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

➤ **Link definition –** If an object has a link, whose minimum limit is **0**, that is, one of the link ends does not necessarily have an object, the new logical object cannot be linked to it (since it may or may not exist in the TQL). For the procedure to define a link condition, see "Defining Link Conditions" on page 72.

For example, in the following scenario you cannot link a logical object to object_15, because it is connected to the host with a minimum limit of **0**:
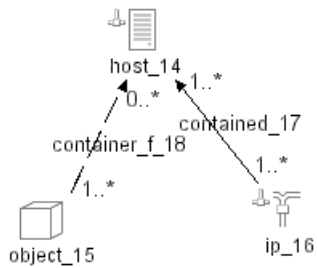
Part III

Views, Organization Rules, and Notifications

6

Introduction to the View Manager

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

This chapter describes:	On page:
Working with View Manager	85
View Manager Shortcut Menus	86
View Manager Tooltips	88

About Creating Managed Views

The View Manager enables you to define managed views, which are a series of rules and definitions for displaying query results. Each view is attached to a specific TQL query (if required, you can define multiple views for each query). By creating managed views, you can define how the query results are displayed in the Map View, including objects, groups and the number of map levels to be displayed.

In addition, by creating view folders, you can determine how the views are arranged in the Map View.

Folders for arranging
views in the map

Explorer pane

Map pane
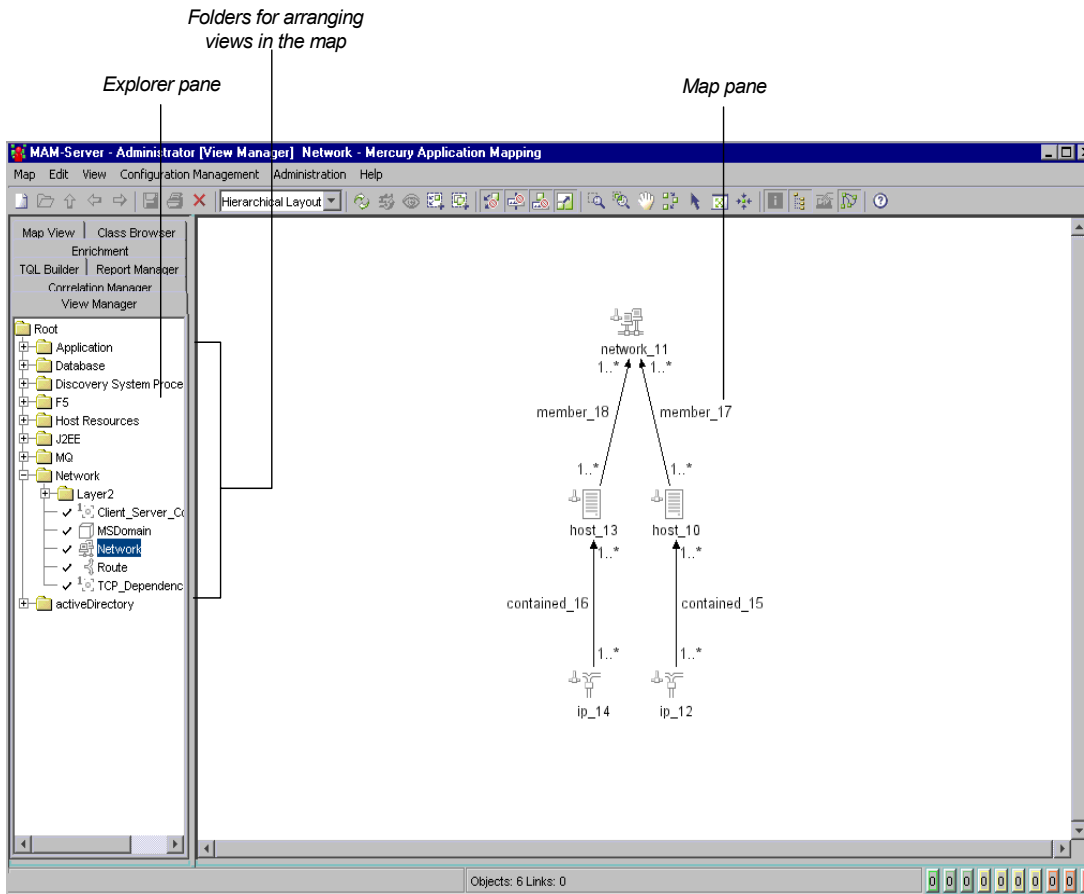
When the View Manager tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **Explorer pane –** 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. If required, you can export views and save them as XML scripts, which can be used for backup purposes. The color of the view icon indicates the view status. For details, see "Creating a View" on page 92.

> ➤ **Map pane –** Displays the currently selected view, which consists of objects that are defined in the TQL query and the relationships between them.

Working with View Manager

Use the View Manager to perform the following functions:

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

➤ Determine how the nodes and links, which are defined in the query, are displayed in the Map View.

➤ Show or hide and delete selected nodes and links in the Map pane.

➤ Define the object's status factor (that is, its system-level significance which affects the status of the object in the Map View.)

➤ Define a grouping option for the nodes displayed in the Map View.

➤ Create and delete tags, which are logical groupings of objects and links, which are defined in the query.

➤ Define the attributes to display in the label of the selected object in the Map View.

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

➤ Define filters for the events to be displayed in real-time in the Map View.

➤ Export and import selected views to and from XML scripts. 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 Map pane area to view the selected layer at different levels of magnification.

View Manager Shortcut Menus

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

➤ "View Folder Shortcut Menu" on page 86

➤ "View Shortcut Menu" on page 86

➤ "Object Shortcut Menu" on page 87

➤ "Link Shortcut Menu" on page 88

View Folder Shortcut Menu

This is the same as the TQL Folder shortcut menu. For details, see "TQL Folder Shortcut Menu" on page 52.

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 Explorer pane:

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 Map View 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 Map View	Enables you to add/remove views to the list in the Map View. When adding a view, a checkmark is displayed beside the selected view in the Explorer pane. This option is especially useful for removing views from the Map View that are created by other users and may not be of interest.
Save	Enables you to save the view to the database. This option is enabled only when a view is created or when changes are made to an existing view.

Option	Description
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.
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.
Notification	Enables you to determine the condition under which notifications are sent.
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 Map View.
Link Rules	Displays the Link Rules dialog box, enabling you to define multiple link rules.

Object Shortcut Menu

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

Option	Description
View Node Definition	Displays the View Node Definition dialog box, enabling you to determine how the nodes defined in the query are displayed in the Map View.
Notification	Displays the Notification dialog box, enabling you to determine the condition under which notifications are sent.

Link Shortcut Menu

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

Option	Description
Straight	(Enabled only for links with angles on the Map pane.) Enables you to straighten links with angles.
Add Rule	Enables you to define the organizational structure of objects displayed in the Map View.
View Link Definition	Displays the View Link Definition dialog box, enabling you to determine how the links defined in the query are displayed in the Map View.

Tip: Right-clicking an empty area in the Map pane 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 15.

View Manager Tooltips

When the cursor is moved over an object in the Map pane, a tooltip is displayed.

The tooltip contains the following information:

Option	Description
TQL	The name of the query.
Node	The name of the node.
Visible	Whether the selected object are visible in the Map View.
Derived	Whether the objects derived from the selected object are displayed in the Map View.

Option	Description
Tags	The tags that have been defined for the selected object and view.
Cardinality	Relevant for links. Cardinality defines how many objects you expect to have at the other end of the link. For example: in a link 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 object.
Status Factor	The relative weight of the selected object, used for determining the status of the object.
Group Object By	The type of group (class, TQL, node, or attribute) the object belongs to, if any.
Minimum in Group	The minimum number of matching objects required to create the group the object belongs to.
Group Label	The label of the group the object belongs to.

7

Defining Managed Views

This chapter explains how to define managed views for viewing the results of TQL queries.

About Defining Managed Views

When creating a new managed view with the View Manager, you must first name the view and attach it to a specific TQL query. You then proceed to define how the results of the query are displayed in the Map View, including groups to regulate the number of objects displayed, and organization rules that determine the number of map levels to be generated.

Managed View Workflow

You create managed views according to the following workflow:

➤ Create a view (for details, see "Creating a View" on page 92).

➤ Define the map categories (for details, see "Defining Map Categories" on page 95).

➤ Define nodes, links and views (for details, see "Defining Nodes, Links and Groups" on page 98).

➤ Add organizational rules to links (for details, see "Adding Organization Rules to Links" on page 109).

➤ Save the managed view (for details, see "Saving the Managed View" on page 112).

Creating a View

The first step when creating a view with the View Manager is to define a unique name and description, attach the view to a specific TQL query and define whether any events related to the view should be displayed in the Map View.

This section includes:

➤ "Merging Identical Instances" on page 93

➤ "Creating a View with the View Manager" on page 94

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 Map View. 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:
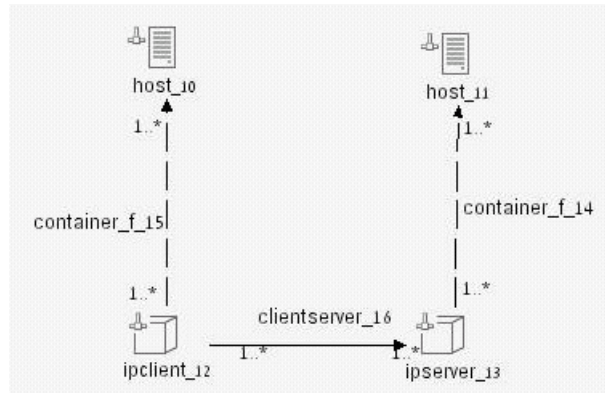
However, if both the server and the client are installed on the same host, one host appears twice.
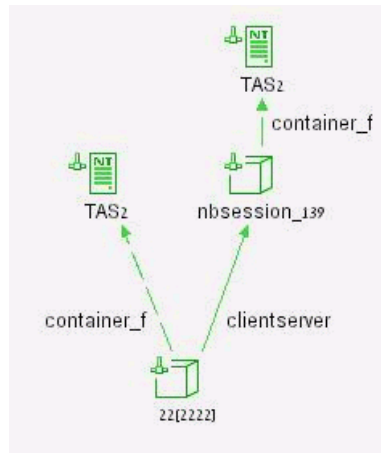
To overcome the redundant host display, you define that identical instances should be merged by selecting **Merge Identical Instances**:

Creating a View with the View Manager

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

To create a view with the View Manager:

1 Click the **View Manager** tab.

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

3 Click the **Properties** tab:

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

➤ From the **Attached TQL** list, choose the TQL to which the view should be attached.

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

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

➤ (Optional) From the **View Icon** list, select a special icon for this view. The view is represented by this icon on the Map View.

Note: If you override the default icon, you cannot see different colors for the priority levels.

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

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

➤ Select **Show Events** to have the system display all events related to the managed view in real time. If **Show Events** is cleared, no events is displayed for this view.

➤ Select **Persistent** to define whether the selected view should 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.

➤ Select **Merge Identical Instances** to automatically display identical instances that appear on the same layer in this view merged into one symbol on the Map View (for details, see "Merging Identical Instances" on page 93). This helps you remove redundant information and sharpen the view's focus.

➤ In the **Send Notification for the following changes** area:

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

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

4 Click the **Advanced** tab to add categories (for details, see below).

5 Click **OK** to display the new view in the Explorer pane.

Defining Map Categories

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 operation category is predefined. To enable a specific map to receive and display additional categories, you add the categories to the **Category Priority** list.

The **Category Priority** list also determines the object color according to the following rules:

➤ Mercury Application Mapping scans the object state for each category 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 category that has the Priority value of 1 has the highest priority, the category that has the Priority value of 2 has the next highest priority, and so forth.

➤ If the object state in the first category (the highest priority value) is not equal to **Normal**, the object 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 387.

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

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

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

➤ In cases where there are two categories 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 object color.

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

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

To add categories:

1 In the View Manager Explorer pane, right-click the object to which you want to add categories and click **Properties** to open the View Properties dialog box.

2 Click the **Advanced** tab.

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

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

4 Click inside the **Category** column and choose the category. The categories that appear in the list are the categories defined in the Category Manager dialog box. For details, see Creating Dynamic Management States, in *Mercury Application Mapping Administration Guide*.

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

Note: You can give the same priority value to two or more categories only if they have the same enumeration definition. For details, see "About Enumerations and Lists" on page 387.

6 Repeat steps 3 to 5 to add another category to the list.

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

Defining Nodes, Links and Groups

After defining the basic parameters of your new view, you can define exactly how the nodes and links of the query are to be displayed in the Map View. You define the significance of the object's status, the grouping options of the nodes (by attributes or by class), the tags of the selected nodes and links, and the labels for each object.

This section includes:

➤ "Defining Nodes or Links" on page 98

➤ "Specifying the Status Factor of an Object" on page 99

➤ "Adding Existing Tags or Creating New Tags" on page 100

➤ "Adding Attributes to the Label of the Selected Object" on page 101

➤ "Grouping Objects" on page 103

Defining Nodes or Links

You can define nodes or links for an object.

To define nodes or links:

In the Map pane, right-click the object (node or link) for which settings should be defined. The View Node Definition or View Link Definition dialog box appears depending on your selection, displaying:

➤ The Status Factor tab (for details, see below).

➤ The Tags tab (for details, see "Adding Existing Tags or Creating New Tags" on page 100).

➤ The Label Format tab (for details, see "Adding Attributes to the Label of the Selected Object" on page 101).

➤ The Group tab (for details, see "Grouping Objects" on page 103).

Specifying the Status Factor of an Object

The status factor of an object is a value (in percent) that defines the significance of the object status. This is useful for reducing the significance of objects whose status can fluctuate on a regular basis, such as processes, which continuously change. By reducing the significance of such objects 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 object status is 2 (5*40/100).

To specify the status factor of an object:

1 In the Map pane, right-click the object (node or link) for which settings should be defined. The View Node Definition or View Link Definition dialog box appears (depending on your selection).

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 object's status. This is useful for reducing the significance of objects whose status can fluctuate on a regular basis, such as processes and CPU. By reducing the significance of such objects 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 object status is **2 (5*40/100)**. For details on status, see "Introduction to the Map View" on page 133

➤ In the **Compound Status Weight** box, enter a value (in percentages) to define the significance of the object's compound status (for details, see compound status on page 494).

4 In the **Propagation** area:

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

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

➤ (Optional) If you selected to propagate the node's status (above), select whether the propagation should be the maximum (**Max**) or average status of the object (**Average**) in the **Propagate Function** list.

Adding Existing Tags or Creating New Tags

Tags are additional organizing filters for presentation purposes. They are logical groupings of objects and links that can be created without any prerequisites (that is, the grouped objects do not need any prior common denominator). You can add existing tags or create new tags and add them to the selected object. The selected tag appears in the object's tooltip.

To add existing tags or to create new tags:

1 In the Map pane, right-click the object (node or link) for which settings should be defined. The View Node Definition or View Link Definition dialog box appears (depending on your selection).

2 Click the **Tags** tab.

3 From the **Available Tags** area, select the tags that you want to add to the selected object's default tags and click the arrow button. The selected tags are displayed in the **Node Tags** area.

> **Tip:** To add all available tags to the selected object's tags, click the double arrow button.
>
> To delete a tag from the View Tags area, select it and click **Delete**. You cannot edit existing tags.

4 (Optional) If required, you can define new tags by entering a name for the tag in the space provided, and clicking **Insert**. The tag that you defined is displayed in the Node Tags area.

5 Click **OK** or click one of the other tabs (for details, see below).

Adding Attributes to the Label of the Selected Object

A label is the title that appears under an object 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.

To add attributes to the label of the selected object:

1 In the Map pane, right-click the object (node or link) for which settings should be defined. The View Node Definition or View Link Definition dialog box appears (depending on your selection).

2 Click the **Label Format** tab.

3 Select the attribute(s) to appear in the label of the selected object (for details, see below).

Formatting Text

+

Select the attribute(s) to appear in the formatted text, by clicking the plus button next to 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 object.
|	Places an **OR** operator between two attributes in the formatted text.
RegExp	Adds a regular expression using regular expression syntax to the formatted text. For examples of how to use regular expression syntax, see "Grouping Attributes" on page 104.

Note: To delete an attribute from the **Format** box, highlight it and click the **Enter** key or the **Delete** key.
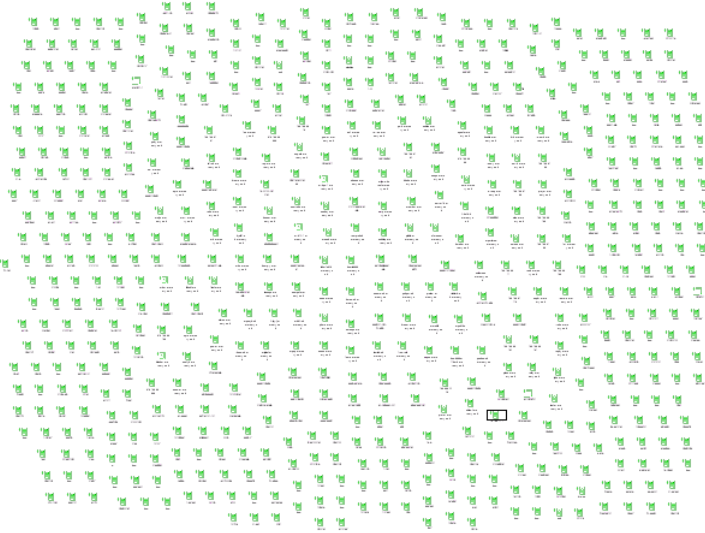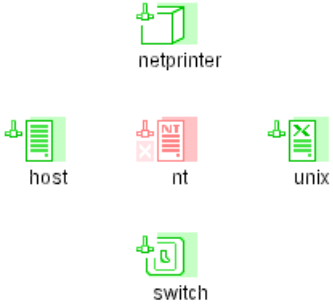
You can define conditions using the following combinations:

➤ AND

➤ OR

➤ AND NOT

➤ OR NOT

Grouping Objects

You can group objects.

The following examples show how the host grouping affects the display in the Map View:

	Non-grouped hosts
	Grouped hosts

Grouping Attributes

You can groups objects from multiple classes in the Map View, 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 objects 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 (an object group inside another object group) to help fine-tune the TQL results that are displayed in the Map View.

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

In the Map View, nested grouping is displayed as follows:

➤ The first layer shows host objects grouped by class:
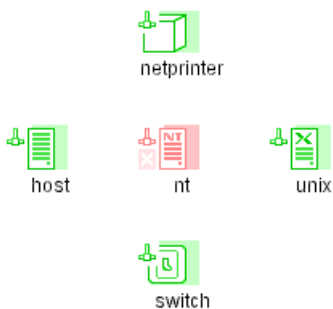
➤ Drill down to the next layer to display host objects grouped by **Oper State**:

➤ Drill down to the next layer to display host objects belonging to a certain Host class and with the same **Oper State**:
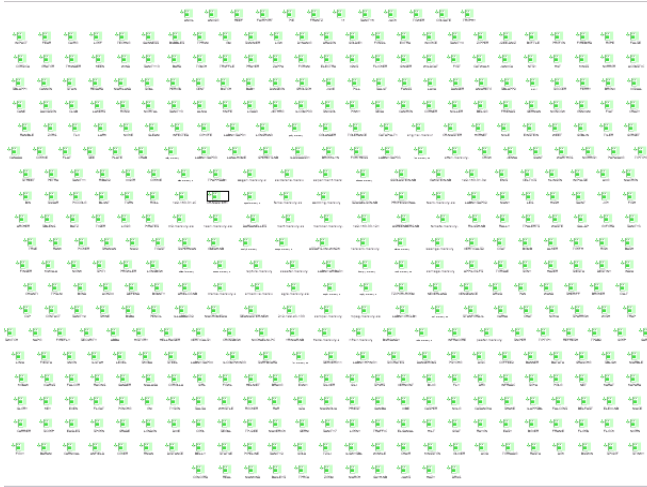
Grouping Objects

You can group objects according to selected criteria.

To group objects:

1 In the Map pane, right-click the object (node or link) for which settings should be defined. The View Node Definition or View Link Definition dialog box appears (depending on your selection).

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

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

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

➤ **Class –** Groups objects in the Map View by class.

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

➤ **Attribute –** Groups objects from multiple classes in the Map View, according to the attribute selected from the list. An object can be a member of multiple groups. The Attribute option enables you to group objects from different classes. The Class and TQL Node options are used to group objects from the same class. For details, see "Grouping Attributes" on page 104).

5 (Optional) In the **Minimum in group** box, enter the minimum number of matching objects required to create the group. By default, there must be at least one object 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, an object group inside another object group. This option provides you with additional fine-tuning of the TQL results that are displayed in the Map View (for more details, see "Nested Groups" on page 105).

8 To nest a group, create one group and return to the Group By tab. 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:

Nested group
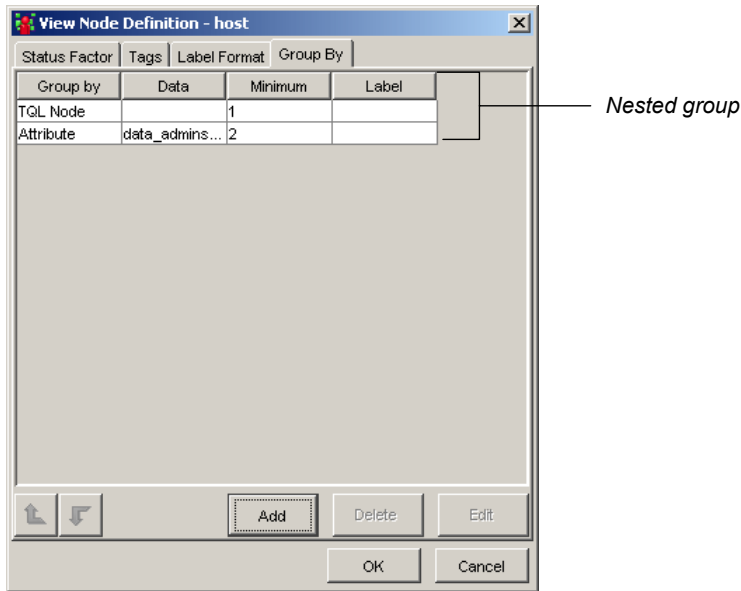
You can repeat this procedure to create as many groups as needed.

Note: After creating several groups, you can change their nesting order by using the arrow buttons at the bottom of the Group By tab.

---

9 Click **OK**. The selected object's definitions are saved.

Adding Organization Rules to Links

After setting the node and link definitions of the view, you can add organization rules to selected links. This enables you to define the organizational structure of objects in the Map View, by displaying selected objects at different levels. When no organization rules are defined, the Map View displays all objects included in the query results on one level by default. For example, if the query results include hosts and networks, both objects are displayed on the same level in the Map View.

To add organization rules:

1 In the Map pane, right-click the link for which you want to define an organization rule.

2 Select **Add Rule**.

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

➤ **Parent –** Select this option to display End1 as the parent of End2. That is, double-clicking the End1 element on one map level displays a new map level with End2 elements.

➤ **Right Sibling –** Select this option to display End1 elements wherever End2 elements are displayed. That is, both objects are displayed on the same map level.

➤ **Left Sibling –** Select this option to display End2 elements wherever End1 elements are displayed. That is, both objects are displayed on the same map level.

➤ **Child –** Select this option to display End1 as the child of End2. That is, double-clicking an End2 element on one map level displays a new map level with End1 elements.

Defining Multiple Link Rules

As an alternative to adding organization rules (for details, see "Adding Organization Rules to Links" on page 109), you can choose to define multiple link rules. You can also define link rules for differentiating between classes that inherit their attributes from one class. For example, if one of the view's nodes is a host, you can specify different links for NT, UNIX, router, and so on. By assigning different organizational rules to different classes, views can encompass and represent additional query data.

Note: If both organization and link rules have been defined, the definitions for the organization rules override the definitions for link rules.

This section includes the following topics:

➤ "Defining Multiple Link Rules" on page 110
➤ "Editing Rules" on page 111
➤ "Deleting Rules" on page 111

Defining Multiple Link Rules

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

To define multiple link rules:

1 In the Explorer pane, right-click the required view for which you want to define link rules.

2 Select **Link Rules** to open the Link Rules dialog box.

3 Click **Add** to add a rule to display the Add Rule dialog box.

4 Select the required End1 object in the **End1** list.

5 Select the required End2 object in the **End2** list.

6 From the **Link** list, select an available link connecting End1 to End2. (Enabled only after you defined both classes.)

7 From the **Rule** list, select an organization rule (**Parent**, **Right Sibling**, **Left Sibling**, or **Child**; for details, see "Adding Organization Rules to Links" on page 109).

8 Click **OK** to save the settings you have defined. You are returned the Link Rules dialog box which now includes the new rule.

9 Repeat steps 3 to 8 to add another rule.

10 Click **OK**.

Editing Rules

You can edit an existing rule.

To edit a rule:

1 In the Link 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 Save the changes.

Deleting Rules

You can delete an existing rule.

To delete a rule:

In the Link Rules dialog box, select the rule you want to delete and click **Delete**.

Saving the Managed View

The last step in defining a managed view is to save it to the database. If you exit the View Manager without saving, a confirmation message is displayed, requesting that you save your managed view.

To save the managed view:

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

2 Click **OK**. The managed view is saved.

To create a view by copying an existing view:

1 Right-click the required view in the Explorer pane and select **Save As**.

2 Enter a new name and description.

3 Click **OK** to save the new view.

Importing a Managed View

You can import XML files that contain saved views to the 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 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 Explorer pane.

8

Sending Notifications

This chapter enables you to notify particular users, group members or unregistered interested parties of changes that occur in views and nodes.

This chapter describes:	On page:
Notification Workflow	114
Adding Availability Rules	115
Defining Unavailability	119
Displaying Availability Results	120
Adding a Status Change Notification to a View or a Node	124
Adding an Event Notification to a View or a Node	126

About Sending Notifications

The Notification sub-system enables you to notify particular users, group members or unregistered interested parties of changes that occur in views and nodes. For each view and node you can define automatic notification, which is sent to a recipient(s) via e-mail or the Mercury Application Mapping messaging system.

These notifications enable you to send information about changes that occur in the availability, status, correlation, or attribute values of managed objects. By using attribute fields that are updated automatically, you can create a notification format that provides up-to-date and comprehensive information on the state of your system.

Notification Workflow

You create notifications according to the following workflow:

➤ Add availability rules (for details, see the next section).

➤ Define unavailability (for details, see "Defining Unavailability" on page 119).

➤ Add notifications for status changes (for details, see "Adding a Status Change Notification to a View or a Node" on page 124).

➤ Add notifications for events (for details, see "Adding an Event Notification to a View or a Node" on page 126).

Adding Availability Rules

After creating a view, you can add availability rules to its nodes and to the view as a whole. These rules enable you to calculate and monitor the availability of the managed objects during predefined periods.

Note: When no availability rules are defined, you cannot track the managed object's availability.

The object's availability is calculated throughout three periods of time:

➤ day

➤ week

➤ month

You can predefine the starting point of each period (for example, the first of each month), or you can start the availability calculation from your current time and date onward. The availability calculation outcome is expressed in percentages that present the relative part of the selected period in which the object is available.

For each calculation cycle, you define one or more thresholds that set the unavailability duration. These thresholds, expressed in percentages, determine for how long the calculation system registers the object as unavailable, before it is considered unavailable for the period of a day, a week or a month. For example, if you set a threshold for a day at 80% and the object is available for only 70% of the time, it is marked as unavailable for the entire day.

The availability calculation is based on the managed object's status – you define from which status level the object is considered as unavailable (for example, Major 8), and Mercury Application Mapping compares the actual object's status with your definition, to decide whether the object is available. If you are interested in the availability of the object's children, you can use the average status of the object's children as the basis for the availability calculation.

Since availability rules are applied to TQL nodes and views, the availability calculations are not limited to one object at a time. You can use different TQL queries and create logical groups of objects to track the availability of network services, which rely on the functionality of interdependent objects.

Once availability calculations are completed, the availability information can be delivered either through e-mail or through an event message. You can define to whom this information is sent and in which format. In each availability notification, a predefined severity level is assigned to the object, which alerts the recipient of the gravity of the object's unavailability. In addition to defining messages, you can generate dashboard indicators, which display the availability results graphically, and System and Map reports, which display the results in a table format.

To add availability rules:

1 Right-click the node or view to which you want to add an availability rule and select **Notification** to open the Notifications dialog box, with the Availability tab open.

2 In the Threshold and Actions section, specify the following availability conditions:

➤ Select **Active** to activate the availability calculation for this condition row. To restore the default values in the **Alert Level** and **Below Uptime** boxes, click the **Reset** button.

➤ To edit the conditions and define actions, select the threshold and click the **Edit** button to display the Availability Definition dialog box.

Note: The **Edit** button is enabled only after you select **Active**.

➤ In the **Below Uptime (%)** box, enter the availability duration, below which the object is considered as unavailable. The applicable range is 0 to 99.

For example: if you specify 75% in the **Below Uptime** box and the object is available for only 70% of the time, it is marked as unavailable.

➤ In the **Alert Level** box, specify which severity level is to be assigned to the object in the availability notifications, once it is considered unavailable.

For example, if you specify **Minor (7)** in the **Alert Level** box, the object severity level in reports, messages and e-mails is **Minor (7)**.

3 In the **Actions** area, click the **Add** button to specify the type and format of the action. The Add New Action dialog box is displayed.

Note: (Optional – if you do not want to define actions for the availability rule, you can still see the availability calculation results through the System and Map Reports).

You can define several actions for each availability rule.

For example, once an object is marked as unavailable, an e-mail could be sent to the system administrator with the label, severity and unavailability duration of the object.

Choose the required action according to the following:

➤ **sendMail** – Send the availability notification in the form of an e-mail to a recipient whose e-mail address you know. For details, see "Sending the Availability Notification as an E-mail" on page 117.

➤ **sendMessageToUser** – Send the availability notification in the form of an event message, to a specific user or a specific group of users. The message is displayed in the recipient's inbox in the Map View. For details, see "Sending the Availability Notification as an Event Message" on page 118.

Sending the Availability Notification as an E-mail

You can send the Availability Notification as an e-mail.

To send the Availability Notification as an e-mail:

1 From the Add New Action dialog box, select **sendMail** and click **Next**, to display the Add New Action dialog box.

2 In the **To** box, enter the full e-mail address of the message's recipient.

3 Click the **Format** button to the right of the **Subject** box to enter a subject for the mail. The Message Format dialog box is displayed (for details, see "Formatting Text" on page 101). The attributes that are displayed are the event attributes. You can send a message that includes not only text but also values that appear in the event. For example: Please note that severity is &activeenvent.eventbase_severity

4 Click **OK**. The defined subject is displayed in the **Subject** box of the Add New Action dialog box.

5 To enter a message in the **Message** box, repeat the above steps.

6 Once you finished entering the required data, click the **Finish** button. You return to the Availability Definition dialog box.

7 If you completed the availability condition definitions, click **OK**.

Sending the Availability Notification as an Event Message

You can send the Availability Notification as an event message.

To send the Availability Notification as an event message:

1 From the Add New Action dialog box, select **sendMessageToUser** to open the Add New Action dialog box.

2 Select the message's recipient from the **User** list.

3 To enter a message in the **Message** box, go to step 3 in "Sending the Availability Notification as an E-mail" on page 117.

Defining Unavailability

You can specify the following unavailability definitions:

➤ from which status the object is considered as unavailable

➤ which status type should be used for the availability calculation

➤ which calculation timetable should be used

To define unavailability:

1 Display the Notification dialog box with the Availability tab open.

2 From the **Not available when status is higher than** list, select the status level from which the object should be considered as unavailable. If a change occurs in the object's status and it becomes higher than the specified status, the system is notified that the object has ceased to be available. The status levels are:

➤ **Critical**

➤ **Major(8)**

➤ **Major(7)**

➤ **Minor(6)**

➤ **Minor(5)**

➤ **Minor(4)**

➤ **Minor(3)**

➤ **Warning(2)**

➤ **Warning(1)**

➤ **Normal**

For example, if you select **Minor (7)** from the **Not available when status is higher than** list, once the object's status reaches **Minor (8)** or higher, it is marked as unavailable.

3 In the **View/Node** area in the **Compared to:** section, select the type of status you want to use for the availability calculation:

- ➤ **Status** – The maximum status of the object and its children.

- ➤ **Compound Status Average** – The average status of the object's children (not including the object's status).

4 In the **View/Node** area in the **Compared to:** section, specify a status type for a group. For details on groups, see "Grouping Objects" on page 103.

5 (Optional) If you want to use a predefined calculation timetable for the availability calculation, select **Calculated according to a timetable**.

6 Once you have entered the availability calculation definitions, click **OK**.

Displaying Availability Results

After you have defined availability rules for views and nodes, you can display and view the availability results through several tools: System Reports, View Indicators, and charts. For detailed step-by-step instructions on how to create and use these tools, see the relevant chapters.

Note: For the following display options, you must define availability rules in advance. Results are calculated only for existing availability rules.

This section includes the following topics:

- ➤ "System Reports" on page 121
- ➤ "View Indicators" on page 121
- ➤ "Charts" on page 122
- ➤ "Setting Up Predefined Timetables for Availability Calculations" on page 123

System Reports

Through system reports you can display the availability of one or more nodes that are part of one view. For details, see Part VIII, "System Reports." Use the **availability** and/or **availabilitylast** functions, when defining the node data to appear in the report:

host_dnsname	availability
10.0.64.110	100%
10.0.64.111	100%
10.0.64.119	100%
10.0.64.120	100%
10.0.64.130	100%
10.0.64.131	100%
10.0.64.197	22%
10.0.64.20	100%
10.0.64.202	100%
10.0.64.21	100%
10.0.64.220	44%
10.0.64.230	100%
10.0.64.254	100%

View Indicators

Through View Indicators, you can display the availability of views in dashboards during a period of day, week or month. For details, see Part V, "Generating View Indicators and Charts."

The following example illustrates a comparison between the availability of two views during a period of one day:
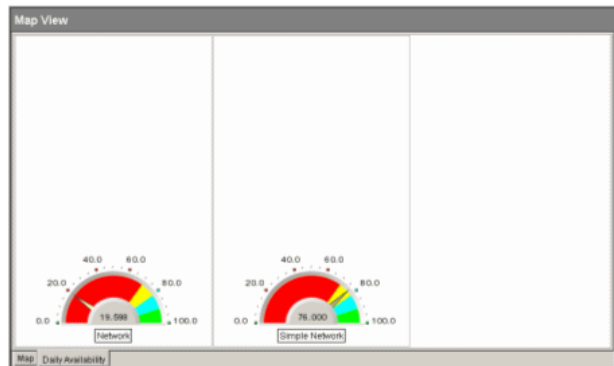
Charts

You can convert the availability information you display in System and Map Reports into charts, and thus view comparisons and trends in the resources and services availability. For details on:

➤ Creating charts, see Chapter 13, "Generating Charts"

➤ Converting System Reports into charts, see "Viewing and Saving the Generated Report" on page 345

➤ Converting Map Reports into charts, see "Converting a Report into a Chart" on page 240

The following example illustrates a comparison between the availability of different network's resources during a period of one day:
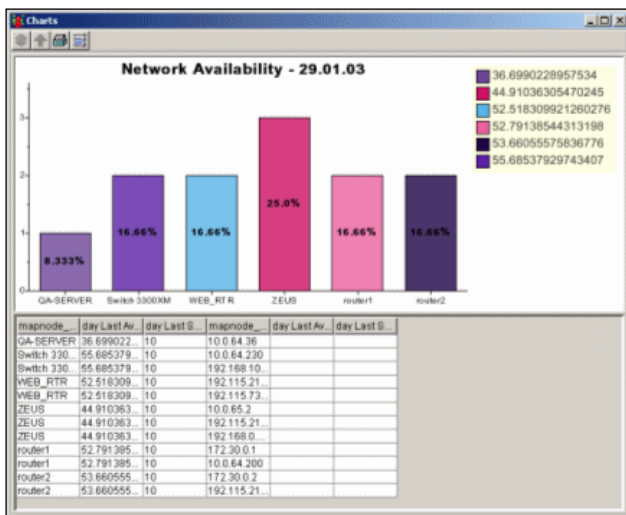
Setting Up Predefined Timetables for Availability Calculations

The timetable definitions are specified in the
**\Mercury\<Mercury Application Mapping root
directory>\root\lib\server\appilogConfig.properties** file, in the Map Server
section:

```
#       Map server

appilog.map.slo.beginOfDay = 08:00:00
appilog.map.slo.beginOfWeek = 2
appilog.map.slo.beginOfMonth = 1

appilog.map.slo.interval.units=minute
appilog.map.slo.interval.basic=1
appilog.map.slo.interval.day=2
appilog.map.slo.interval.week=10
appilog.map.slo.interval.month=24
```

*Calculating
starting point*

*Calculation
intervals*

➤ The first three definitions, which include the words **beginOf**, specify the
starting point of each availability calculation: an hour in the day, a day
in the week, and a day in the month.

For example, in the above definitions the availability calculation starts at
8:00 AM for a day, on a Monday for a week, and on the first of the month
for a month.

➤ The next five definitions, which include the word **interval**, specify the
intervals between each calculation in each of the three calculation cycles.

For example, in the above definitions, the basic unit of the intervals is
one minute. The calculations are performed in the following intervals:
for a day – every two minutes, for a week – every ten minutes, for a
month – every twenty-four minutes.

Note: If you changed the **appilogConfig.properties** file, you need to restart
the server for the changes to take effect.

Adding a Status Change Notification to a View or a Node

Status change notification enables you to notify particular users, group members or unregistered interested parties of changes occurring in the Map Status of views and nodes. By defining notifications that are sent automatically once a status changes, you can get constant updates, via Mercury Application Mapping inbox or an e-mail, on the state of your IT infrastructure and business services. Thus, you can handle these changes after their occurrence, without the need to continuously keep track of each view and node separately.

When defining a status notification, you set a status range for each node or view. If a status falls into the defined range, a status notification is sent. For example, if you defined a status range between Major(8) and Critical for a certain node, once the node's status reaches Major(8), a status notification is sent automatically to specified people.

To add a status change notification to a view or a node:

1 In the View Manager, right-click the view or node for which you want to define a status notification and select **Notification** to open the Notification dialog box with the Availability tab open.

2 Select the **Status** tab.

3 Click **Add** to define the status notification for the selected view or node to open the Definition dialog box.

4 From the **From** list, select the Map Status that is to be the lowest end of the change range. When the view/node status reaches this level, a notification is sent about the change occurrence.

For example, if you select Minor(5), when the status of the view/node reaches this level, a status change notification is sent. If the status exceeds this level, a notification may or may not be sent, depending on the upper end of the range you define in the **To** list.

5 From the **To** list, select the Map Status that is to be the upper end of the change range. When the view/node status goes beyond this level, this notification is not sent.

Note: You cannot select a status level in the **To** list that is lower than the one you selected in the **From** list. However, you can select from both lists the same status level, if you want to define a status notification for a specific status level.

For example, if you select Minor(7) when the status of the view/node falls between this level and level you define in the **From** list, a status change notification is sent. If the status is higher than Minor(7), a notification is not sent.

6 Click **Add** to display the Add New Action dialog box.

Select one of the following actions:

➤ **sendMail –** For sending an e-mail status notification, to a specific user or a specific group of users. This e-mail is automatically sent to the defined recipient(s), once the status change occurs. For details on sending notifications by e-mail, see "Sending the Availability Notification as an E-mail" on page 117.

➤ **sendMessageToUser –** To send a status notification to be displayed in the recipient's inbox in the Map View. This status notification is not registered in the Event tabs and browsers. For details on sending events to users and groups, see "Sending the Availability Notification as an Event Message" on page 118.

Note: You can define several actions for each status notification.

7 Click **Next**.

8 (Optional) To add a further action to the status notification, click **Add** in the Definition dialog box. Then, repeat step 6.

9 Click **OK** in the Definition dialog box.

The Notification dialog box is displayed, and the new status notification is displayed in the list.

10 (Optional) To define additional status notifications for the selected view/node, click **Add** and repeat steps 3 to 9.

11 Click **OK** to save the status notification.

Adding an Event Notification to a View or a Node

You can automatically notify users of active events that occur with regard to objects that appear in views and nodes. For each view and node, you can define the kind of event for which you want to be notified, the notification method and the recipients of the event notification. You or other interested parties can receive constant updates, via Mercury Application Mapping inbox or an e-mail, on significant events that occur in your IT infrastructure and business services.

This section includes the following topics:

➤ "Adding an Event Notification to a View or a Node" on page 126

➤ "Editing an Event Notification" on page 128

➤ "Removing an Event Notification" on page 129

➤ "Editing an Action" on page 129

➤ "Removing an Action" on page 130

Adding an Event Notification to a View or a Node

You can add an event notification to a view or a node.

To add an event notification to a view or a node:

1 In the View Manager, right-click the view or node for which you want to define an event notification.

2 Select **Notification** from the menu to open the Notification with the Availability tab open.

3 Select the **Event** tab.

4 Click **Add** to define the event notification for the selected view or node to open the Event Rule dialog box.

5 In the Conditions area, click **Add** to open the Condition dialog box.

The Conditions area is used to define the conditions of the rule that determines the event type.

6 To define a new rule condition for the event, do the following:

➤ From the **Attribute** list, select the event's attribute to be included in the condition.

Tip: You can use the Event tabs and browsers to accurately define the event type of which you want to be notified. For details, see Chapter 12, "Event Management."

➤ From the **Operator** list, select the required operator (for details, see "Operator List" on page 45 and "Defining Attribute Conditions for Nodes" on page 71).

➤ From the **Value** area, select or enter a value to be included in the condition, according to the event's attribute you selected.

➤ Click **OK**. The condition is added to the table at the top of the Event Rule dialog box.

Tip: You can create complex rules by defining and combining several conditions for one rule.

---

7 The Actions area is used to define the actions to be performed when the conditions of the rule are met, that is, when an event that meets your defined condition(s) occurs. To define the notification's format and recipients, click **Add** to open the Add New Action dialog box.

8 In the Available Actions area, select one of the following actions:

➤ **sendMail –** For sending an e-mail status notification, to a specific user or a specific group of users. This e-mail is automatically sent to the defined recipient(s), once the status change occurs. For details on sending

notifications by e-mail, see "Sending the Availability Notification as an E-mail" on page 117.

➤ **sendMessageToUser –** To send a status notification to be displayed in the recipient's inbox in the Map View. This status notification is not registered in the Event tabs and browsers. For details on sending events to users and groups, see "Sending the Availability Notification as an Event Message" on page 118.

Note: You can define several actions for each event notification.

---

9 Click **Next**.

10 To add a further action to the event notification you defined, click **Add** in the Definition dialog box. Then, repeat step 7.

11 Click **OK** in the Event Rule dialog box. The Notification dialog box appears, and the new event notification you defined is displayed in the list.

12 (Optional) If you want to define additional event notifications for the selected view/node, click **Add** and repeat steps 4 through 11.

13 Click **OK** to save the event notification you defined.

Editing an Event Notification

You can edit an existing event notification.

To edit an event notification:

1 In the View Manager, right-click the view or node for which you want to edit the event notification.

2 Select **Notification** from the menu to open the Notification with the Availability tab open.

3 Select the **Event** tab.

4 To edit a condition, select it in the **Conditions** area, and click **Edit** to open the Condition dialog box.

5 Go to step 5 for more details.

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

Removing an Event Notification

You can remove an existing event notification.

To remove an event notification:

1 In the View Manager, right-click the view or node for which you want to remove the event notification.

2 Select **Notification** from the menu to open the Notification with the Availability tab open.

3 Select the **Event** tab.

4 To remove a condition, select it in the **Conditions** area, and click **Delete**.

5 Click **OK** to save the change.

Editing an Action

You can edit an existing action.

To edit an action:

1 In the View Manager, right-click the view or node for which you want to edit the action.

2 Select **Notification** from the menu to open the Notification with the Availability tab open.

3 Select the **Event** tab.

4 To edit an action, select it in the **Actions** area, and click **Edit** to open the Send Mail dialog box.

5 Go to step 8 for more details.

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

Removing an Action

You can remove an existing action.

To remove an action:

1 In the View Manager, right-click the view or node for which you want to remove the event notification.

2 Select **Notification** from the menu to open the Notification with the Availability tab open.

3 Select the **Event** tab.

4 To remove an action, select it in the **Actions** area and, click **Delete**.

5 Click **OK** to save the change.

Part IV

Viewing Multi-level Maps with the Map View

9

Introduction to the Map View

This chapter introduces the Map View, which enables you to view multi-level maps displaying the results of a TQL query or an instance view. In addition, you can manage the events occurring in the system, at the object 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:
Working with Map View	135
Understanding Map View Concepts	136
Map View Shortcut Menus	138

About Introduction to the Map View

The Map View is displayed by selecting the Map View tab in the Explorer pane, and enables you to view the results of your TQL queries based on the view defined with the View Manager or the instance views created in Map View. For information on how to build a view from a TQL query, see Chapter 4, "Defining Topology Query Language (TQL) Queries". For information on how to create an instance view, see "Defining Instance Views" on page 205.

The maps displayed in the Map View can be displayed in multiple layers, depending on whether organization rules are defined in the view, as described in Part III, "Views, Organization Rules, and Notifications." The Map View also enables you to view all events associated with the objects displayed in each map.
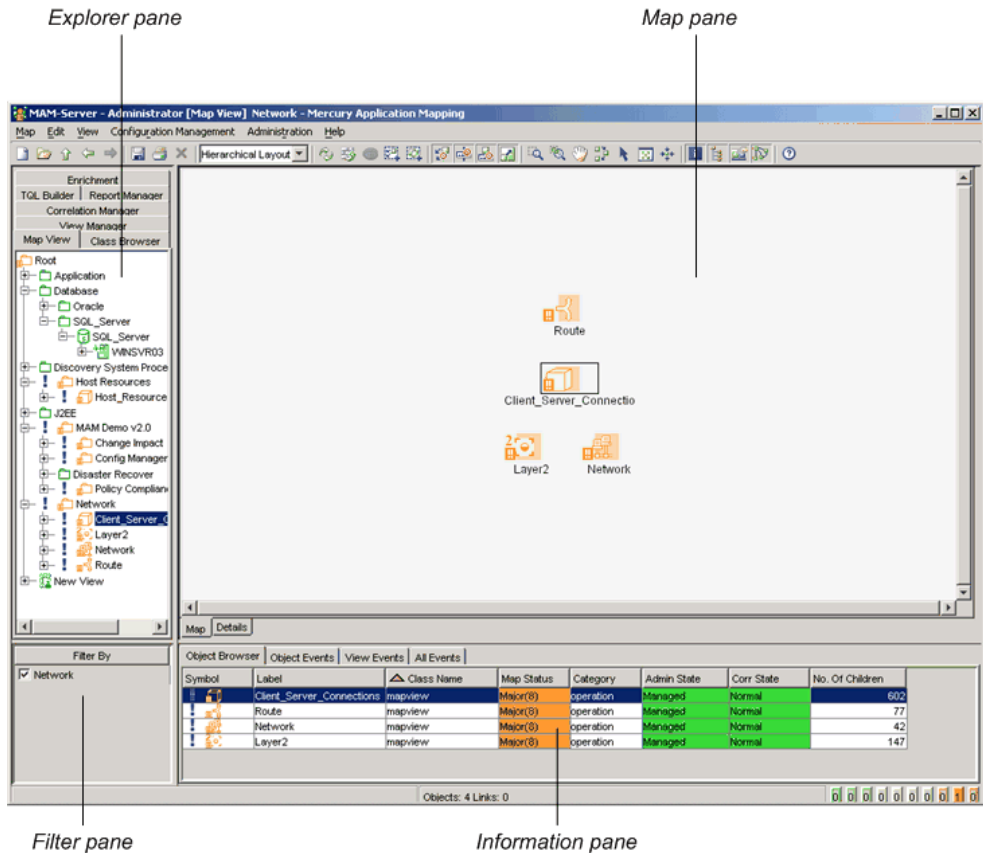
When the Map View tab is selected, the Mercury Application Mapping main window is divided, as follows:

➤ **Explorer pane –** Displays a hierarchical tree structure of the map objects based on a selected view. You can navigate between map layers by either selecting layers from the Explorer pane or by drilling down. For example, by selecting the view of a network in the Explorer pane, you can drill down to view the objects belonging to the selected network at the host level.

Selecting an object in the Explorer pane simultaneously selects it in both the Map and Information panes. The color of the selected view's icon is determined by its severity level.

➤ **Filter pane** – Displays the tags defined in the View Manager for the objects currently displayed in the Map pane. Tags are logical groupings of objects. You can show or hide these tags, as required.

➤ **Map pane** – Displays the results of the view currently selected in the Explorer pane, and consists of objects 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, Circular, Symmetric, Manual, Orthogonal, or Tree). For details, see "Improving Views" on page 172.

➤ **Information pane** – Displays all events that are related to the objects in the selected view. This pane includes the following tabs: Object Browser, Object Events, View Events, All Event, and Inbox. For details, see "The Event Tabs in the Information Pane" on page 218.

Working with Map View

Use the Map View to perform the following functions:

➤ Display the views based on a TQL query created with the View Manager or instance views created in Map View.

➤ Create a new instance view.

➤ Show or hide selected objects, links or views.

➤ Edit labels of selected objects in the Map pane.

➤ Display the neighbors of a selected object, which are connected to the object using the various link types.

➤ Display and edit the object's attributes, according to your access rights.

➤ View the number of events per severity level or per view (displayed in the status bar).

➤ Acknowledge events relating to the selected objects and to the objects displayed in the levels below the selected objects (the compound objects).

➤ Show or hide tags, which represent logical groupings of the objects and links that are defined in the query.

➤ Display the selected object's events and compound events.

➤ Zoom in and out of the 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 Map View Concepts

The objects and links that are represented in the Map View have unique topology data in addition to their basic attributes. This topology data is based on, and produced according to, the way these objects and links 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:

➤ "Map View Tooltip" on page 136

➤ "Object Browser Tab" on page 137

Map View Tooltip

When you point to an object in the Map pane, a tooltip is displayed that contains the object statuses in the selected categories of the view:

The tooltip contains the following information:

➤ **Label** – The label of the selected object, as defined in the Label Format tab of the Node/Link Definition dialog box. For details, see "Adding Attributes to the Label of the Selected Object" on page 101.

➤ **Class** – The class of the selected object as defined in the class model.

➤ **Category** – The category name.

➤ **Status** – The severity level of the category.

➤ **Corr** – Indicates the severity of the object that is affected by a correlation rule (in the corresponding category).

➤ **Is new** – Indicates that the object has a new event that is neither acknowledged nor suppressed.

➤ **Comp** – Exists in parent objects only and indicates the maximum status of the selected object's children.

For example, if the status of one of the child objects is **5** and the status of another child is **3**, then the compound status is **5**.

Object Browser Tab

The Object Browser tab lists all the objects of the selected view that are currently displayed in the Map pane. (To display the Object Browser tab, select **View > Information Browsers > Object Browser**.) In addition to the information displayed in the tooltip, the Object Browser contains the following information:

➤ **Symbol** – The icon defined for the class of the selected object.

➤ **No. of Children** – The number of children and links 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 an object in the Object Browser tab also selects it in both the Map and the Explorer panes.

Map View Shortcut Menus

The Map View contains different shortcut menus, depending on your selection in the Map View.

This section includes:

➤ "View Shortcut Menu" on page 138

➤ "Object Shortcut Menu" on page 140

➤ "Multiple Objects Shortcut Menu" on page 145

➤ "Link Shortcut Menu" on page 147

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 Explorer, 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.
Rebuild	Refreshes the objects and links displayed in the selected view.

Option	Description
Snapshot	➤ **Save** – To take a snapshot of a view – for details, see "Taking a Snapshot of a View" on page 394 • **View snapshot** – 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 – for details, see "Comparing Snapshots Taken at Different Times" on page 396
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).
Reports	**Asset Report** – creates a report that lists all the objects in a selected view and their attribute values – for details, see "Generating an Asset Report" on page 264 **Host Dependency Report** – create a report that lists all the links 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 – for details, see "Generating a Dependency Report" on page 266 **Event Report** – generates a report displaying the events for the objects in the selected view – for details, see "Generating an Event Report" on page 269

Option	Description
Events	(These options are enabled for parent objects only.) Select one of the following options: ➤ **Compound Event Browser –** Displays the Compound Event Browser dialog box that enables you to view active events related to objects in lower maps that are connected to the selected view. Through this dialog box you can perform other actions. ➤ **Compound Event Log –** Displays the Compound Event Log dialog box that enables you to view all the raw events related to objects in lower maps that are connected to the selected view. Through this dialog box you can also create a chart.
Gold Master Report	Generates a report that compares the configuration of a Gold Master object to other objects of the same class – for details, see "Gold Master Comparison" on page 411

Object Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a selected object in the Map, Explorer or Information panes.

Option	Description
Hide	Enables you to hide the selected object. By hiding one or more objects, you can focus on the objects 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 Map View.
Note	Creates personal notes that include comments, and attaches them to nodes or events in the Map View – for details, see "Adding Notes to Displayed Objects" on page 175

Option	Description
Edit Label	Enables you to edit the label of the selected object, which is defined in the View Manager as the object's name.
Delete	Enables you to delete the selected object from the view and from the Mercury Universal CMDB.
	Note: When you delete a parent object, the selected object and its children are removed from the view, but only the selected object is removed from the database.
Get Neighbors	Enables you to display the interdependencies of a selected object (that is, its neighbors) in different contexts: a specific layer, a specific view or the entire CMDB. For details on the Get Neighbors options, see "Displaying Interdependent Objects" on page 181.
Host Reports	(Enabled only for objects that belong to the host class.) Displays the Reports on Host dialog box, enabling you to display system reports that relate to the data of the selected host.
Show Object Attributes	Displays the Object Attributes dialog box, enabling you to view and edit object attributes, according to your access rights.
Insert Object	Displays the Insert Object dialog box, enabling you to manually add new objects to the database and to the view. For details, see the *Mercury Application Mapping Administration Guide*.

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 objects that are affected by the selected object. If only one correlation rule is defined for this object, 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 object(s) that affected the selected object. If only one correlation rule is defined for this object, 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 object. ➤ **Compound Event Browser** – (This option is enabled for parent objects only.) Displays the Event Browser dialog box, enabling you to view all the events related to the selected object and to objects in lower map levels that are connected to the selected object. ➤ **Event Log** – Displays the raw event history of the selected object.
Host Reports	(Enabled only for objects that belong to the host class.) Displays the Reports on Host dialog box, enabling you to display system reports that relate to the data of the selected host.

Option	Description
Acknowledge	(This option is enabled only when unacknowledged events exist for the selected object or its children.) Enables you to acknowledge events, as follows: ➤ **Object –** Acknowledge all the events relating to the selected object. (This option is enabled only when unacknowledged events exist for the selected object itself.) ➤ **Compound –** Acknowledge all the events related to objects in lower map levels and to the object itself (useful when viewing the top map level).

Option	Description
Actions	Select one of the following options: ➤ **Set Admin State –** Enables you to manually change the Admin State of an object, so that it is no longer a managed object in the system. For example, by setting the Admin State of a router to **Testing**, the router has no influence on other objects in the system, and the other objects 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 object, according to your access rights. For details on defining the event category, see "Sending an Event" on page 234. ➤ **Clear Events –** Enables you to clear the object's state and return it to **Normal**. For details on clearing the object's state, see "Clearing All Events" on page 234. ➤ **Add Discovery Pattern –** Displays the Add Discovery Pattern dialog box, enabling you to manually invoke a Discovery Pattern for the selected object. You can use this option to discover additional information about the object 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 object.)
HostPing	(This option is enabled for host and IP objects only.) Enables you to ping the selected object.
Gold Master Report	Generates a report that compares the configuration of a Gold Master object to other objects of the same class – for details, see "Gold Master Comparison" on page 411

Multiple Objects Shortcut Menu

The table below contains a brief description of each option in the shortcut menu, displayed by selecting several objects and right-clicking them in the Map pane (for details, see "Selecting Multiple Objects" on page 44).

Note: Use Windows conventions—SHIFT+ARROW key, CTRL+ARROW key—to select adjacent (for the Explorer tree only) or nonadjacent objects.

You cannot select multiple objects from different view layers.
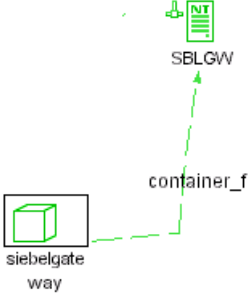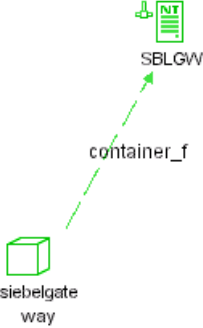
Option	Description
Hide	Enables you to hide the selected object. By hiding several objects, you can focus on the object(s) that are of interest to you. This option is useful for display purposes and can assist you in creating reports and diagrams, which are based on the current display of the Map View.
Delete	Enables you to delete the selected object(s) from the view and from the CMDB. **Note:** When you delete a parent object, the selected object and its children are removed from the view, but only the selected object is removed from the database.
Show Path	(Enabled for two selected objects only.) Enables you to show the objects that compose the path between two selected objects. This option can help you detect connection problems between objects.
Insert Link	(Enabled for two selected objects only) Display the Insert Link dialog box, enabling administrators to manually add links between objects that are saved to the database. Adding a link to the database triggers the Discovery Methods to begin collecting the required information regarding the new link.

Option	Description
Actions	**Add Discovery Pattern –** Displays the Add Discovery Pattern dialog box, enabling you to manually invoke a Discovery Pattern for the selected object. You can use this option to discover additional information about the object 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.
Gold Master Report	Generates a report that compares the configuration of a Gold Master object to other objects of the same class – for details, see "Gold Master Comparison" on page 411

Link Shortcut Menu

This section contains a brief description of each option in the shortcut menu displayed by right-clicking a link in the Map pane:

Option	Description
Hide	Enables you to hide the selected link.
Straight	(Enabled only for links with angles on the Map pane.) Enables you to straighten links with angles, such as the one that appears in the following example: To straighten the link, right-click the blue square in the middle of the angle and select **Straight**. The link is straightened, as seen below:
Delete	Enables you to delete the selected link from the view and from the database.

Option	Description
Show Object Attributes	Displays the Object Attributes dialog box, enabling you to view and edit link attributes, according to your access rights. For details, see "Viewing and Editing Object Attributes" on page 176.
Map Reports	Display the Map Reports dialog box, enabling you to create and display Map Reports that relates to the selected objects data. For details, see Part V, "Generating View Indicators and Charts."
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 objects that are affected by the selected link. If only one correlation rule is defined for this link, the Impact window is directly displayed. For further details, see "Showing a Correlation Impact" on page 189. ➤ **Show Root-cause** – 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 object(s) that affected the selected link. If only one correlation rule is defined for this object, the Correlation window is directly displayed. For further details, see "Showing Root Cause Objects" on page 191.
Report	➤ **Event Browser** – Displays the Event Browser dialog box, enabling you to view the events related to the selected link. For details, see "Event Browser" on page 227.

Option	Description
Actions	Select one of the following options: ➤ **Set Admin State** – Enables you to manually change the Admin State of a link, so that it is no longer a managed object in the system. For example, by setting the Admin State of a router to Testing, the router has no influence on other objects in the system, and the other objects 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 link's state, according to your access rights. ➤ **Clear All Events** – Enables you to clear the link's Operational State and return it to Normal.
Gold Master Report	Generates a report that compares the configuration of a Gold Master object to other objects of the same class – for details, see "Gold Master Comparison" on page 411

10

Using Map View

Map View displays the results of a pattern or an instance view. This chapter explains the different functionalities of Map View, including how to navigate between different map levels, define the layout of view's layers and display the interdependencies of a selected object.

This chapter describes:	On page:
Understanding Map View	152
Navigating in Maps	152
Understanding Layer Layouts	154
Improving Views	172
Viewing and Editing Object Attributes	176
Viewing Layer Details	178
Viewing Layers with a Large Number of Objects	180
Displaying Interdependent Objects	181
Showing a Correlation Impact	189
Showing Root Cause Objects	191
Using Meta Links	194
Saving a Map View	196
Printing the Map View	198
Sorting, Hiding, and Displaying Columns	200
Searching for Objects	202

Understanding Map View

Map View displays the results of an instance view created in the Map View or a view based on a TQL query created with View Manager. For information on how to create an instance view, see Chapter 11, "Defining Instance Views". For information on how to create a view with View Manager, see Chapter 7, "Defining Managed Views".

Instance views and views based on a TQL query are represented by different icons in View Explorer, as described in the following table.

Icon	What it represents
	A view based on a TQL query.
	An instance view

Map View shows the results of the currently selected view in the View Explorer, and consists of objects and links that are defined in the view. You can change the way the maps are displayed by selecting one of the view layouts (Hierarchical, Symmetric, Manual, or Orthogonal). For details, see "Understanding Layer Layouts" on page 154. You can also modify the values of the default layer layouts, as described in "Defining a Layer Layout" on page 173.

Navigating in Maps

The Map View enables you to navigate between different map levels (from the top level down to the object level) to focus on specific areas of the network, according to your requirements. In addition, the Map View enables you to view the attributes of any displayed object and add a background graphic to the map.

When a view is selected in the Explorer pane, the map for that view is displayed in the Map pane. The icon displayed for each object represented in the map is determined by its class, as defined in the class model, and includes the label defined for that class in the View Manager.

If the object has a shadow, you can double-click the object to drill down to a lower map level, created by a child organizational rule set up in the View Manager.

Host_Resources

For example, if the IP object is defined as a child of the host object in the View Manager, you can double-click any host icon in the Map View to display a map of IPs connected with that host.

To drill down from a selected parent object, you can also use the Open option on the **Map** menu, or the **Open** button on the toolbar.

Right-click an object and select from the shortcut menu one of the **Get Neighbors** options, to display a map containing the interdependencies of the selected object in different contexts: a specific layer, a specific view or the entire CMDB. Selecting the **Get Neighbors > Layer** option is useful when working with maps that display a great number of objects, 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 object and other objects in your system, as described in "Displaying Interdependent Objects" on page 181.

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 links 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 links that are displayed in the Map view. 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 links 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.

Bounding rectangle of the expanded node

Bounding rectangle of the nested graph

Frame bounding rectangle of the nested graph

View rectangle

Margin spacing

Always zero

Nested view spacing
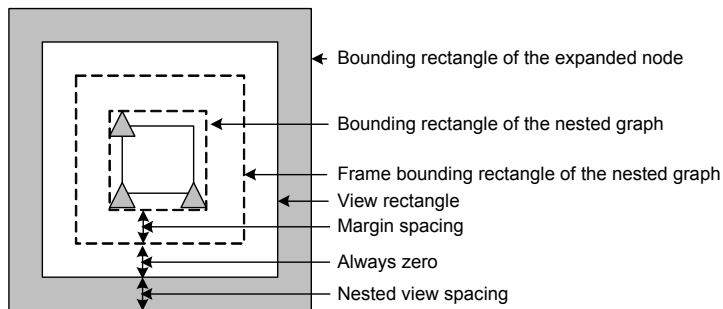
➤ **Intergraph Edge Spacing** – If intergraph edges are routed orthogonally, you can specify horizontal and vertical spacing for the edges.

➤ **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
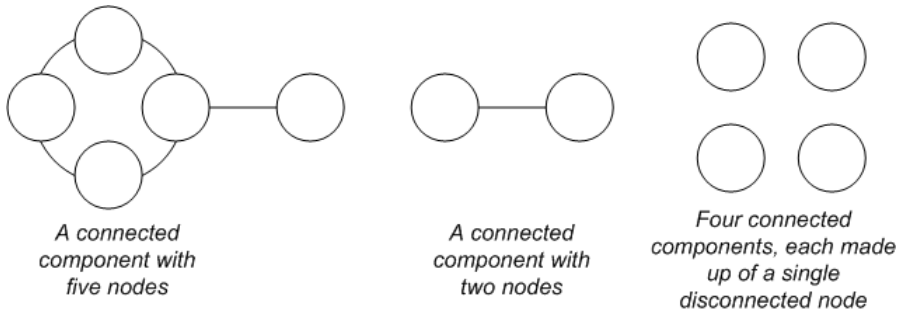
➤ **Components** – The spacing between the components in a disconnected graph consists of both a constant value and a proportional value based on the sizes of the components.

➤ **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.
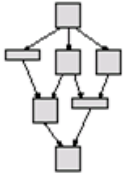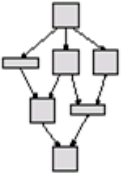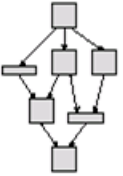
There are several layout options you can apply to a Hierarchical layout, as described below.

➤ **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).	

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).	
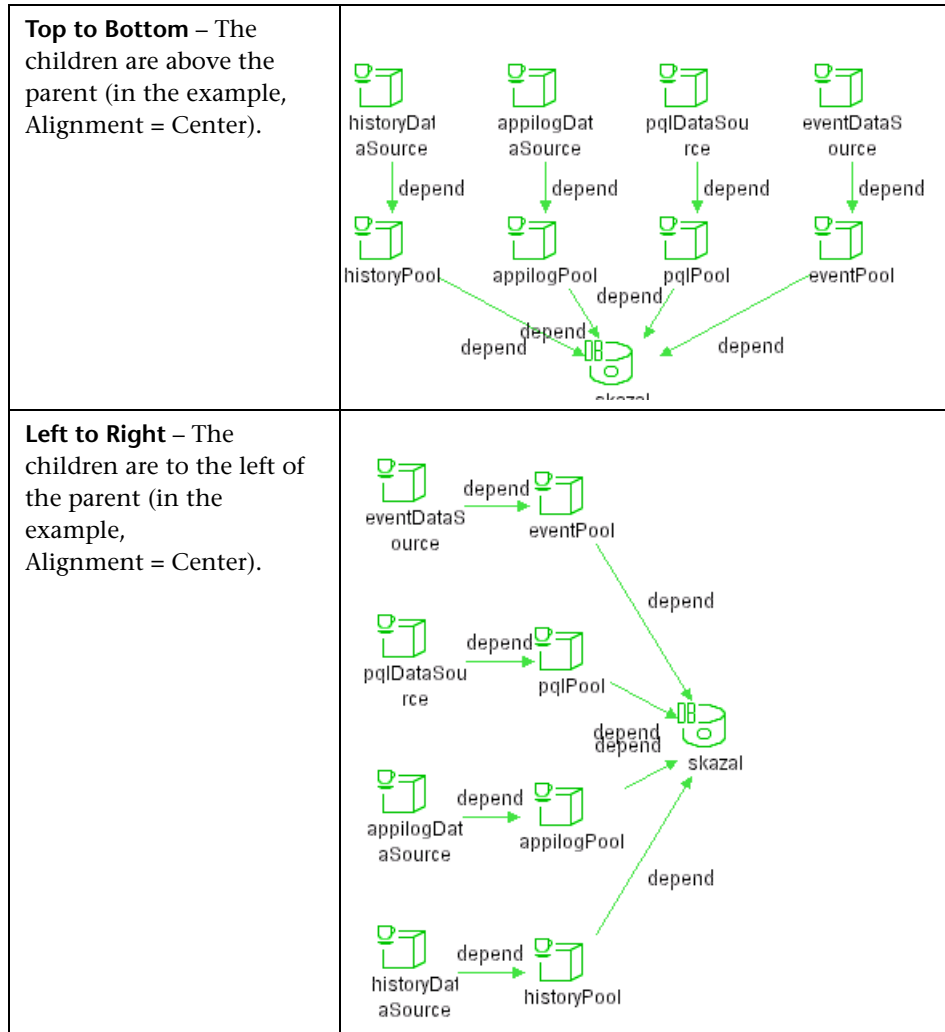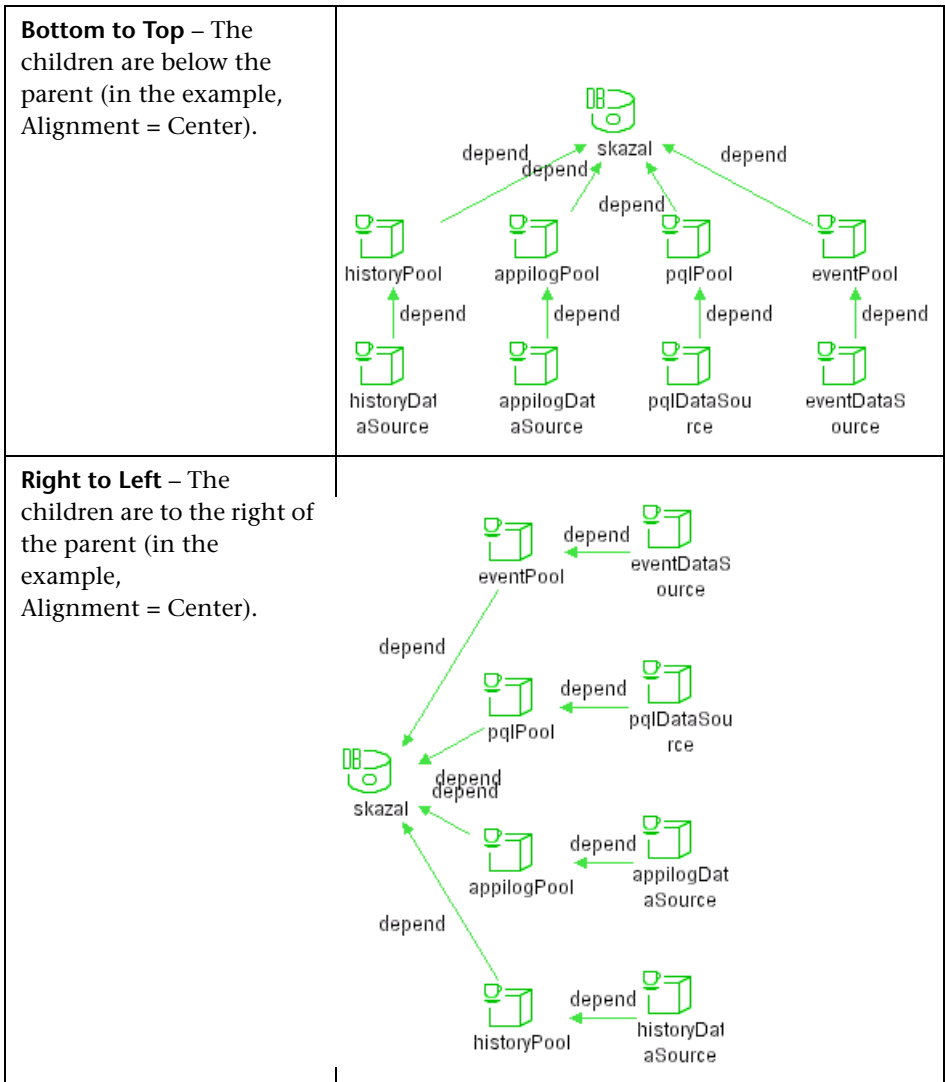
➤ **Level Alignment** – This option allows you to control the way nodes are aligned on a vertical level.

| **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.
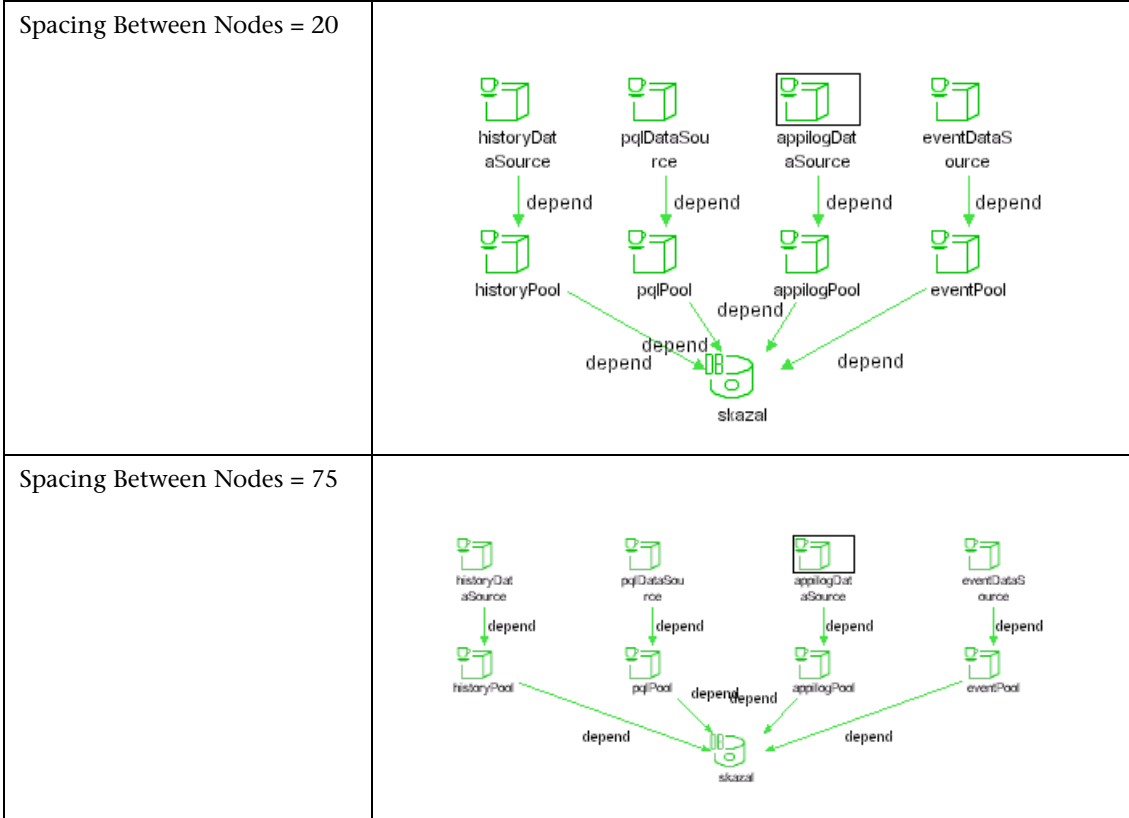
Spacing Between Nodes = 20	
Spacing Between Nodes = 75	

➤ **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	
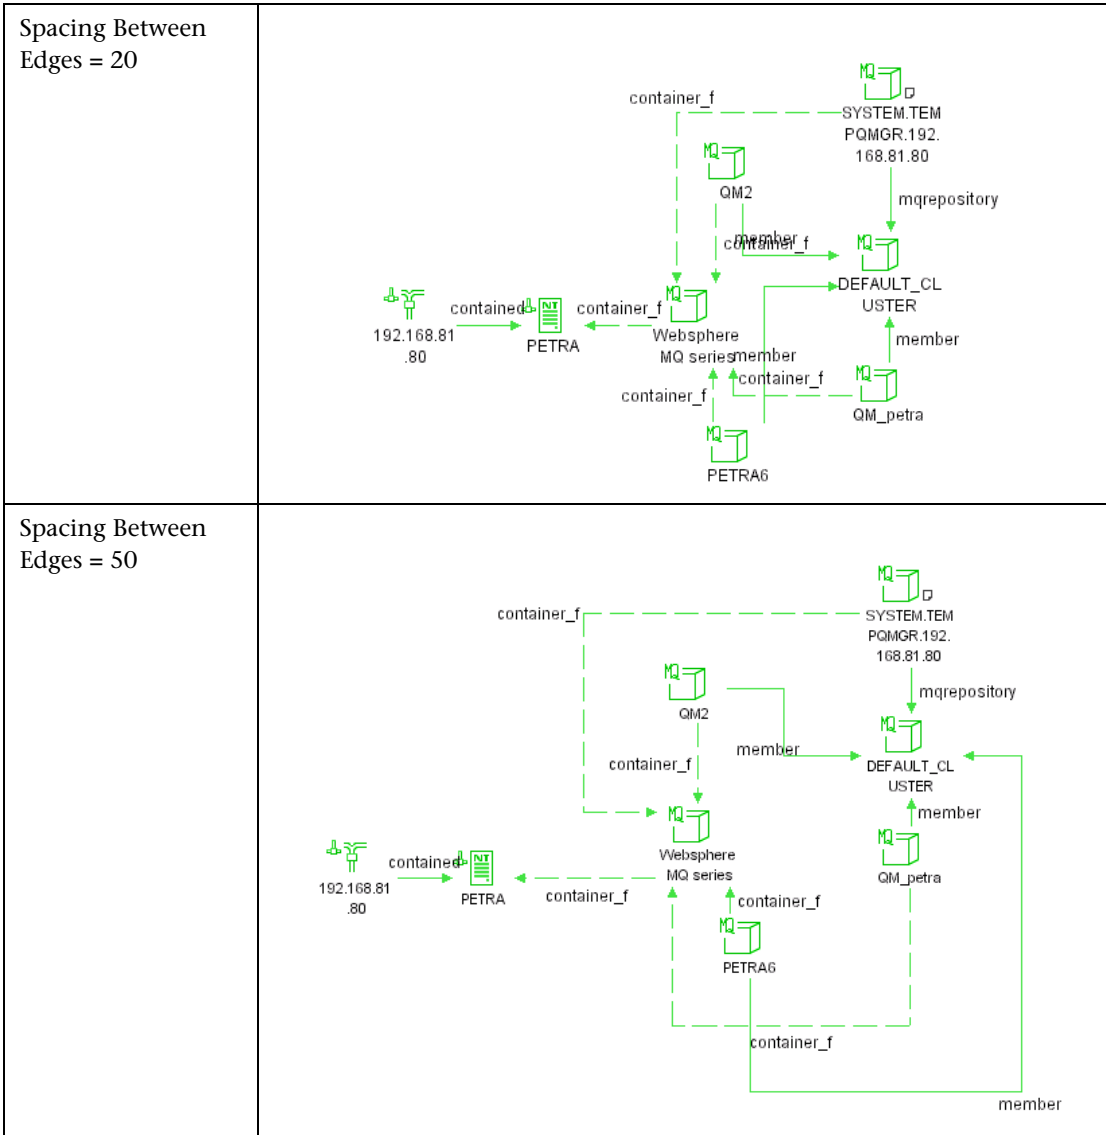
➤ **Vertical Spacing –** Sets the minimum vertical distance between two neighboring nodes on different levels.

➤ **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 link between objects. This option is useful when layers have nodes with a very large number of connecting links.

Orthogonal – To have orthogonal links between objects.	
Polyline – To have straight lines between objects.	
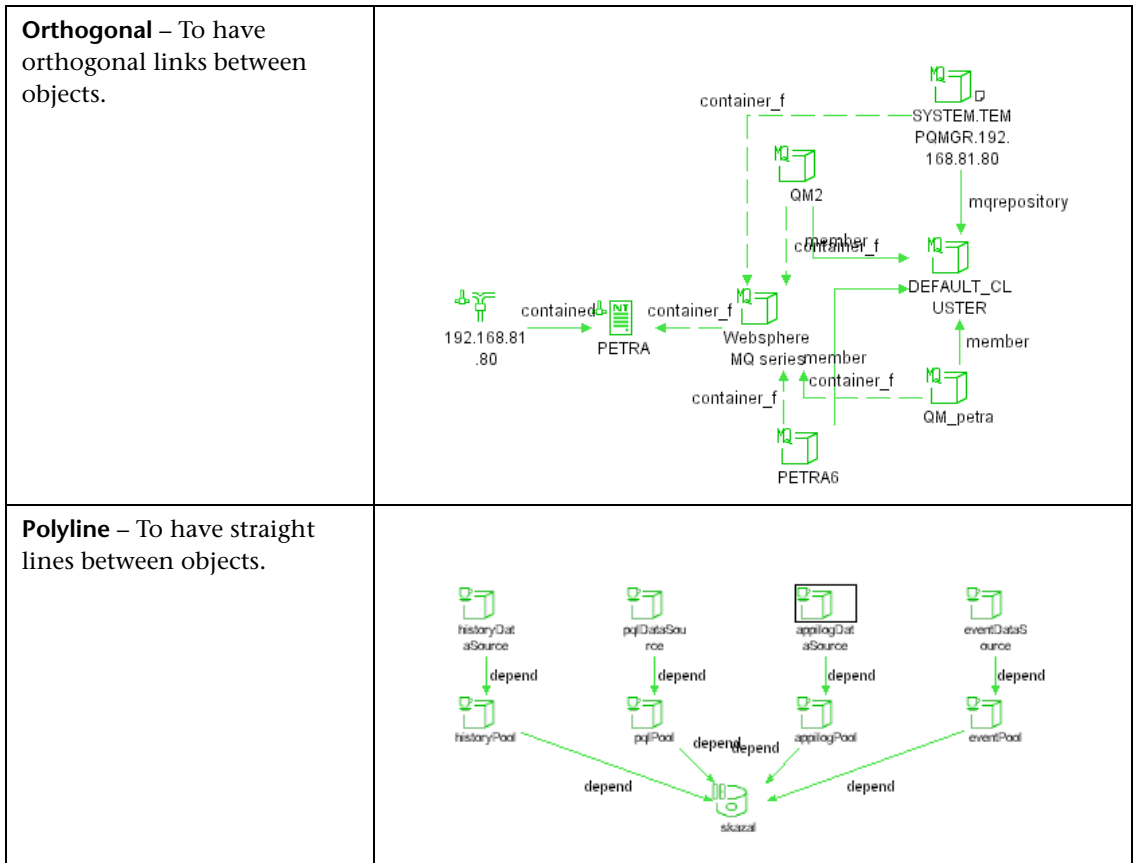
➤ **Orthogonal Routing** – Routes edges horizontally and vertically.

➤ **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.
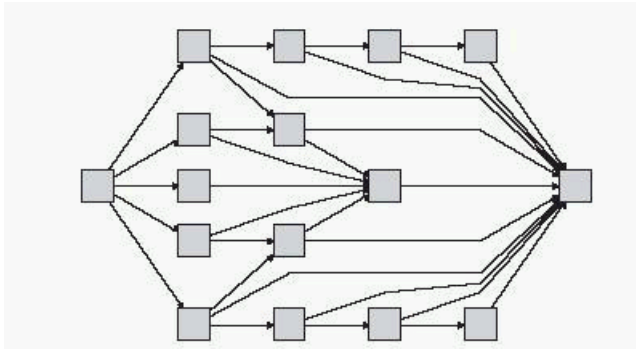
➤ **Variable Level Spacing** – Adjusts the spacing between pairs of neighboring levels according to the density of edges between the levels.

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.

➤ **Undirect Layout** – Edge direction is not used to build the levels of the hierarchical drawing.

Orthogonal Layout

The orthogonal layout routes links horizontally and vertically. This results in links bending at ninety-degree angles only.

➤ **Horizontal Spacing –** To set the minimum horizontal distance between objects.
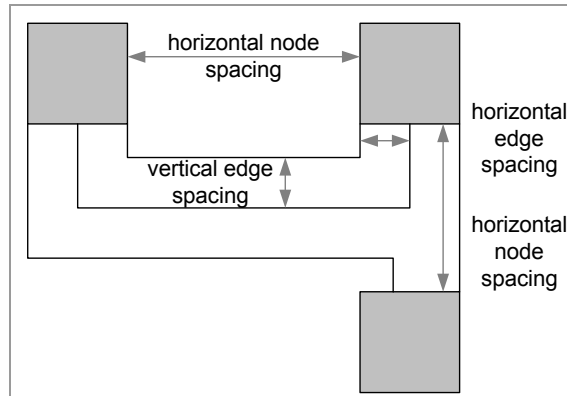
➤ **Spacing Between Nodes** – Sets a space around each node in the layer. The default value is 20.

➤ **Spacing Between Edges** – Sets a space between elements in the layer.

➤ **Vertical Spacing –** To set the minimum vertical distance between objects.

➤ **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 Size** – For future use.

Symmetric Layout

The symmetric layout emphasizes the symmetries that may occur in a graph.
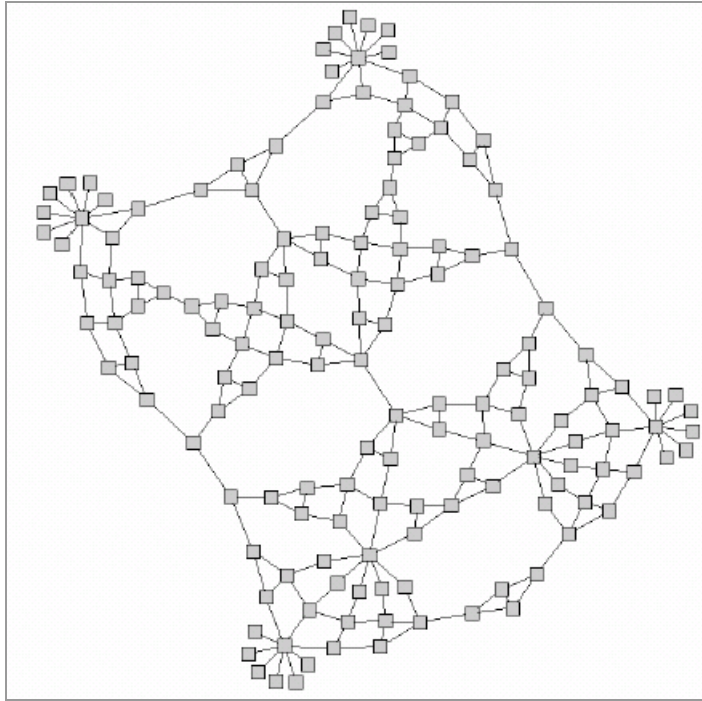
This section describes the symmetric layouts that are available.

➤ **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.
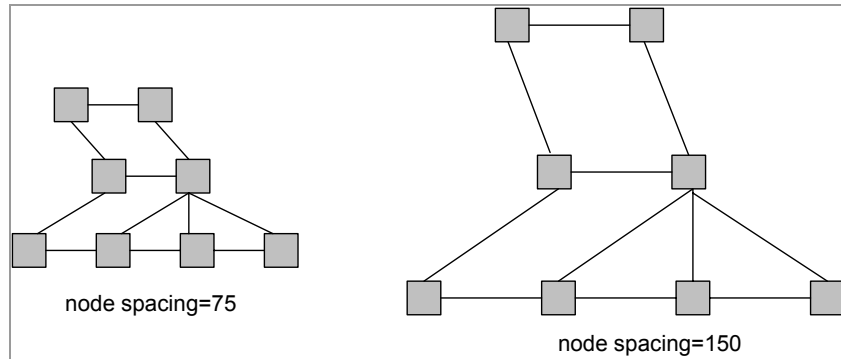
node spacing=75

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
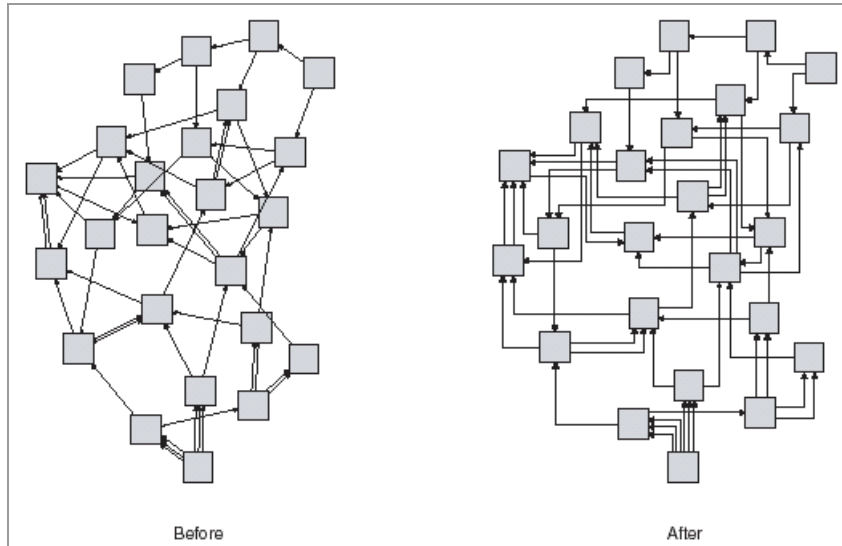
This section describes the routing layouts that are available.

➤ **Horizontal Spacing** – To set the horizontal space around each node or each object.

➤ **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	
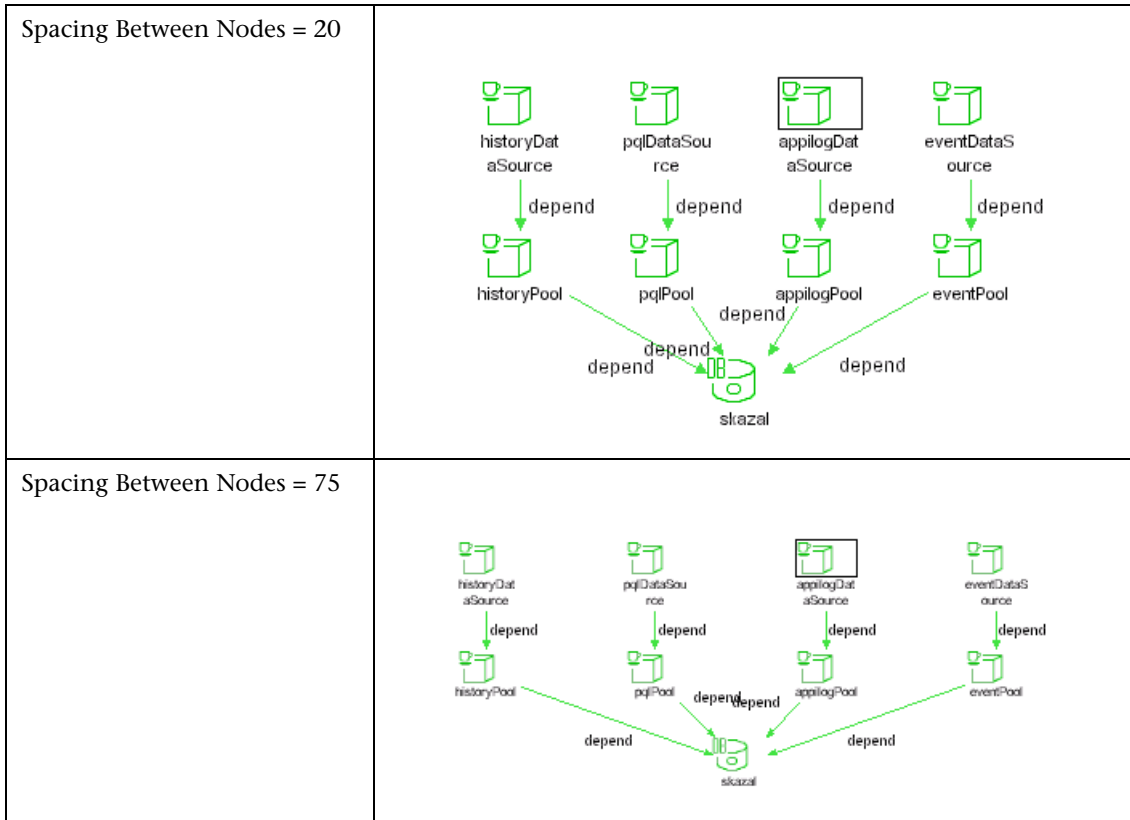
➤ **Spacing Between Edges** – Sets the horizontal space between any two neighboring parallel edges. Edge spacing can also affect the size of a node to which many edges are attached, unless the node is set to remain fixed at its current size. The settings for edge spacing are constant values.

For example:

Spacing Between Edges = 20	
Spacing Between Edges = 50	
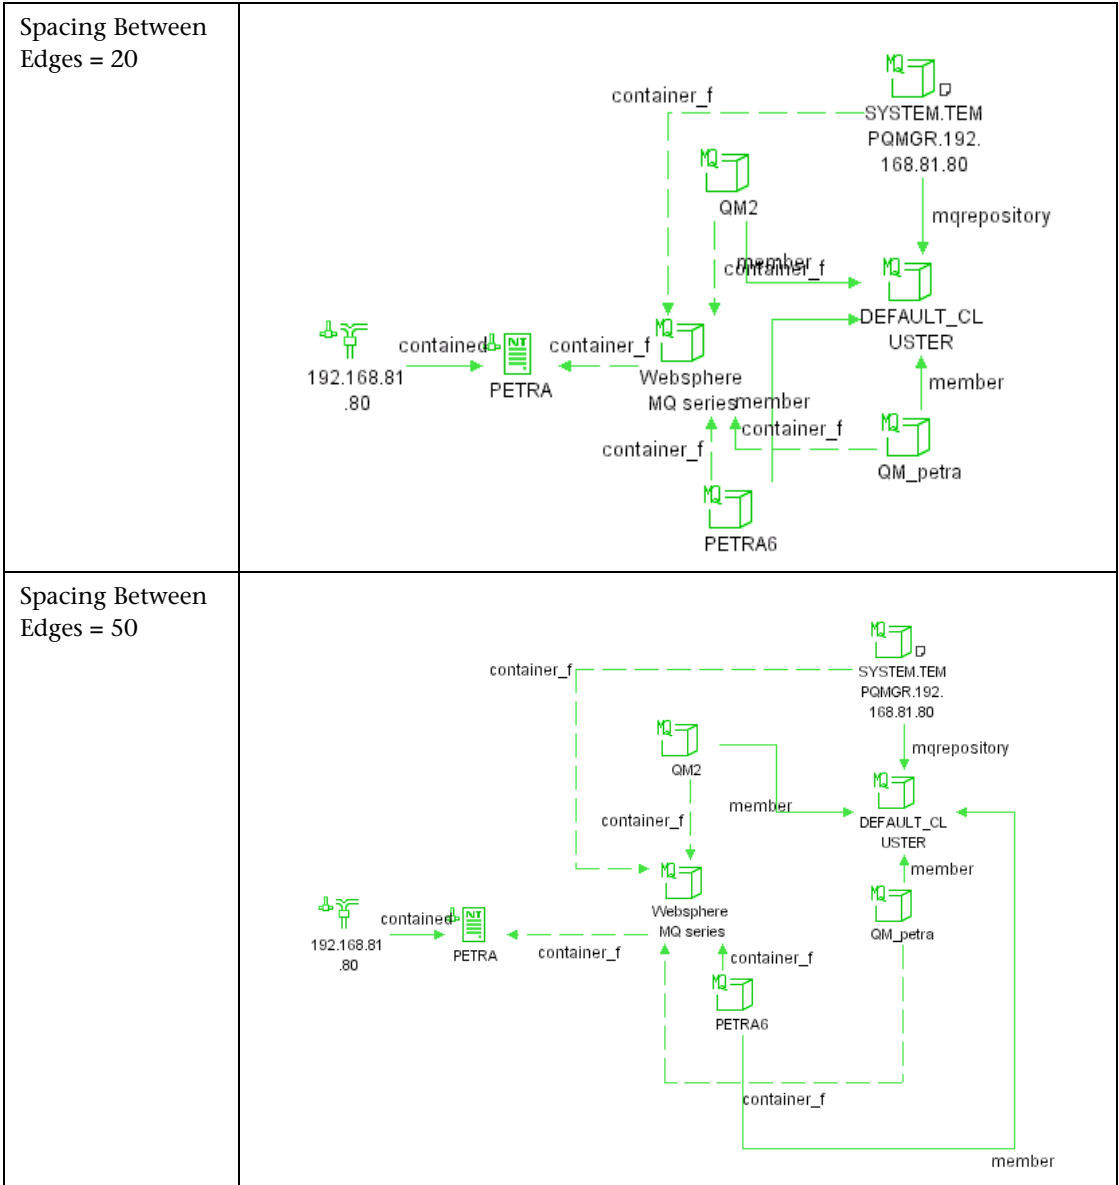
➤ **Vertical Spacing –** To set the vertical space around each node or each object.

   ➤ **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 link spacing when more than one link 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.

Improving Views

You can improve a view by:

➤ Modifying the default layer setup to achieve the clearest presentation possible

➤ Resizing the labels of all objects that are displayed in one layer in the Map View

➤ 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 Map View.

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 155.

➤ **Disconnected** – Defines general layout parameters for disconnected nodes and connected components. For details, see "Disconnected Layout" on page 156.

➤ **Hierarchical** – Shows the precedence relationships of the nodes in the selected layer. For details, see "Hierarchical Layout" on page 157.

➤ **Orthogonal** – Routes links only horizontally and vertically. For details, see "Orthogonal Layout" on page 165.

➤ **Symmetric** – Displays a clear representation of complex networks. For details, see "Symmetric Layout" on page 167.

➤ **Routing** – Leave nodes where they are in the graphic representation and reroutes the edges orthogonally. For details, see "Routing" on page 169.

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 Map View

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 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 Map pane.

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 Objects

You can create personal notes that include comments, and attach them to nodes or events in the Map View. 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 an object.

To add a note to an object:

1 From the Explorer, Map, or Information pane, right-click the object 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

Objects 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 object or right-click the object 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 object and select **Note** > **Delete Note**.

Viewing and Editing Object Attributes

The Attributes dialog box enables you to view and edit specific object attributes, depending on your user access rights to the system (for details on access rights, see the *Mercury Application Mapping Administration Guide*).

To view and edit object attributes:

1 In the Explorer pane, Map pane or Information pane, right-click an object and select **Show Object Attributes** to open the Object Attributes dialog box.

The Object Attributes dialog box enables you to view and edit the attributes defined for the selected object. The name of the selected object appears at the top. When the **Locked on object** check box is checked, Mercury Application Mapping displays only the attributes of the object in the box at the top. Clear the **Locked on object** check box to display the attributes of any object you select.

Note: Only the attributes that were marked as **Editable** and **Visible** in the Class Browser appear in the Object Attributes dialog box. For details, see "Creating Classes" on page 374.

The Object Attributes dialog box has the following tabs:

➤ **Alphabetic** – Displays the object's attributes in alphabetical order.

➤ **Categorized** – Displays the object's attributes according to the class to which it belongs. Click the **Expand** button to display the object attributes for that class.

➤ **Object Attributes History** – Displays a list of attributes in which changes have occurred.

Note: The Object Attributes History tab only displays the attributes whose values have been defined as **Change Monitored** in the Add Attributes dialog box in the Class Browser. For details, see "Creating Classes" on page 374.

The Object 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 Objects Attributes History tab by selecting an attribute from the **Filter by** list.

2 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.

3 Make the required changes in the dialog box that opens and click **OK**.

Note: After you have made your changes, you may have to wait a few minutes until the change appears in the Object Attributes History tab since the Object Attributes History tab is only updated every ten minutes.

4 Click **Apply** in the Object Attributes dialog box to save your changes or click **OK** to save your changes and close the Object Attributes dialog box.

Viewing Layer Details

The Details option enables you to display the attributes of all the objects that are located in the same view layer in one table. The presented details are divided according to the object classes, as is shown in the following example:

To display a layer's details:

1 Select the layer whose details you want to display on the Explorer pane.

2 On the **View** menu, click the **Details** tab to open the Details table on the Map View.

To display object attributes from a different class:

From the class list on the left, select another class name. The attribute list on the right changes according to the selected class.

To update the layer's details table:

Press the **F5** key or move to another Map View tab and return to the Details tab.

To filter the layer's details table:

1 Right-click the details table, and select **Filter** to open the Filter dialog box.
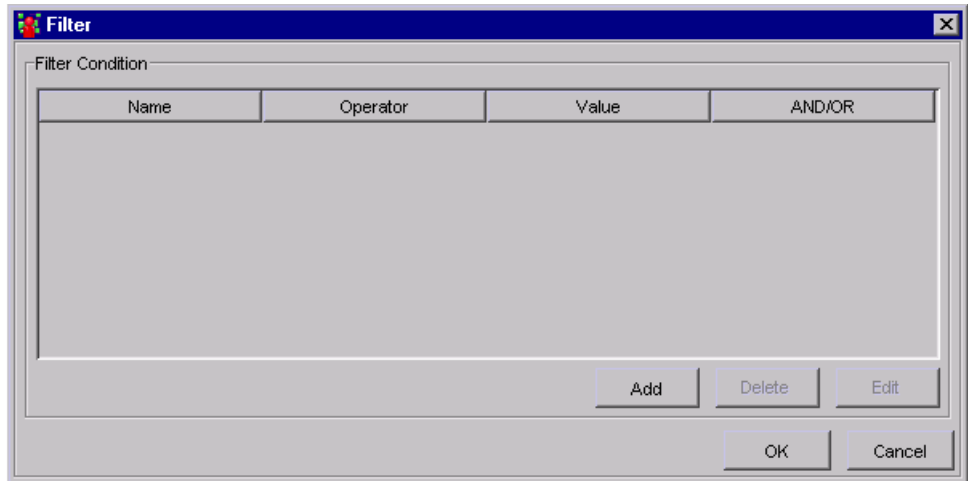
2 To create a new filter, click **Add**, and follow the steps described in "Defining an Event Filter" on page 230.

3 To apply a filter to the details tables, open the Filter dialog box, select the required filter and click **OK**.

4 To remove a filter, right-click the details table and select **Clear Filter**.

Note: The Details filters can be reused during one session only.

To close the Details tab:

Right click the **Details** tab, and select **Remove Details** or select **Details** again from the **View** menu.

Viewing Layers with a Large Number of Objects

A layer cannot display more than 1000 objects in one map view. You can, however, display layers that contain more than 1000 objects and less than 3001 in a table format. When using this option, the layer's objects 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 objects, 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 objects among several layers. For using organization rules, see the previous section.

To display a layer that has more than 1000 objects:

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 objects of the layer:

The arrows in the lower-right corner of the Details window enable you to display the other objects that are in the layer in quantities of 1000 objects.

Note: You can configure the maximum number of objects that can be displayed in a table format (default is 3000), and the number of objects that are displayed in each bulk in the Details window (default is 1000). For details on configuring these parameters, see the *Mercury Application Mapping Administration Guide*.

Displaying Interdependent Objects

The **Get Neighbors** option enables you to display the interdependencies of a selected object in different contexts: a specific layer, a specific view, or the entire CMDB. By using this option, all objects that are connected to the selected object in a chosen context are displayed, either on the Map pane or in a separate window. This way, you can visualize, locate and manage the interdependencies that exist among your IT resources.

The **Get Neighbors** option opens a submenu, which contains three main commands—one for each interdependency context—Layer, View, and Database.

Selecting one of the **Get Neighbors** commands opens an interdependency map for the selected object. Subsequently, Mercury Application Mapping provides you with additional commands, which enable you to expand the interdependency data that the map displays to other connected objects, 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**, **Next on View**, or **Next on Database**. These commands are enabled only after using the Layer, View, or Database options.

By applying one of these commands to a connected object, additional interdependency information is displayed. You can then continue using this command, and display interdependency data for any displayed object.

To display interdependent objects:

1 In the Map, Explorer or Information pane, right-click the object whose interdependent objects you want to display.

2 Click **Get Neighbors** and then select one of the following commands:

➤ **Layer –** Displays all objects that are connected to the selected object via links in a specific layer. This option actually removes from the displayed layer all objects that are not connected to the selected object, thus enables you to focus on the interdependencies of a specific object 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 object is displayed on a regular view's layer:
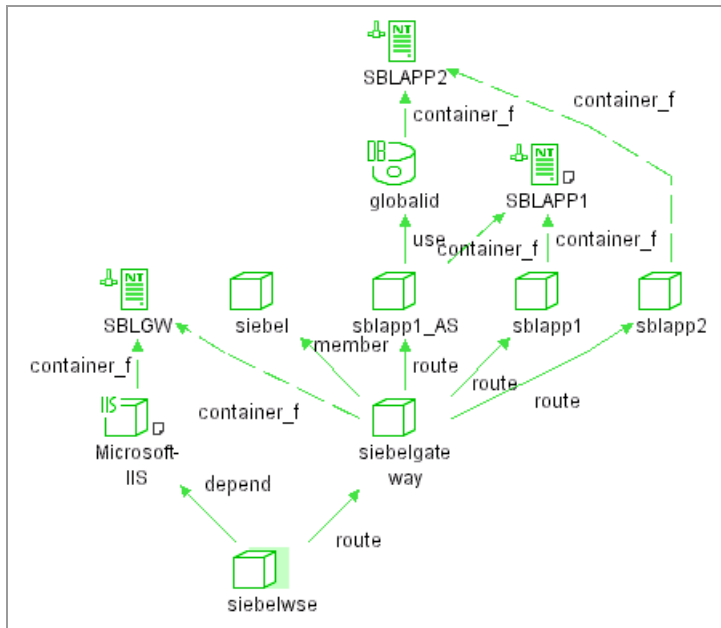
When using the **Layer** option, Mercury Application Mapping hides all objects that are not connected to siebelgateway on this layer, and displays only its interdependent objects:
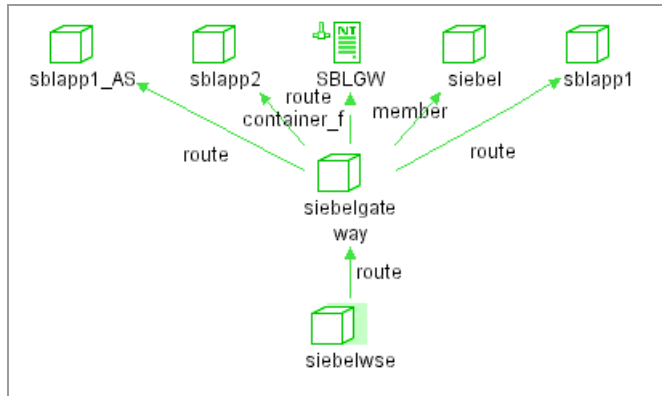
➤ **View** – Displays all objects, connected to a selected object, that are part of the view's TQL, regardless of the layer in which they are located according to the view's organizational rules. The View's content is displayed in a new window.

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 objects that are connected to router1 in a given view, the following window is displayed:
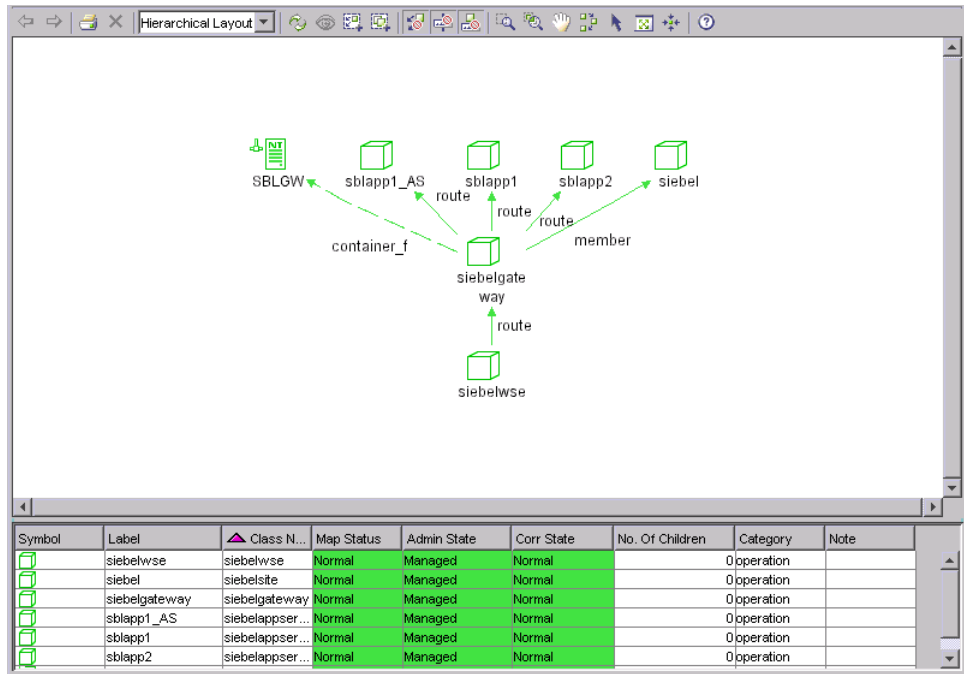
As you can see in the above window, there is one additional object connected to siebelgateway, in comparison to the Layer display. This object is located in a lower layer level, and therefore it does not appear on the Layer display.

Database – Displays objects connected to the selected object from the CMDB. When you select the **Database** option, the Advanced Get Neighbors dialog box opens from which you can either choose to display all the related objects from the CMDB or only related objects from specific classes.

- Select **Get All Objects from DB** to display all related objects from the CMDB.

- Select **Filter Objects From DB** to display only related objects from specific classes. When you select **Filter Objects from DB**, the Advanced Get Neighbors dialog box has the following fields:

Field	Description
Classes	The name of the class.
Count of Actual Objects	The number of related objects found in the CMDB of that class type.

Select the the check boxes of the classes whose objects you want to be displayed. 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 objects that are connected to siebelgateway, regardless of their position in the existing TQL scenarios:
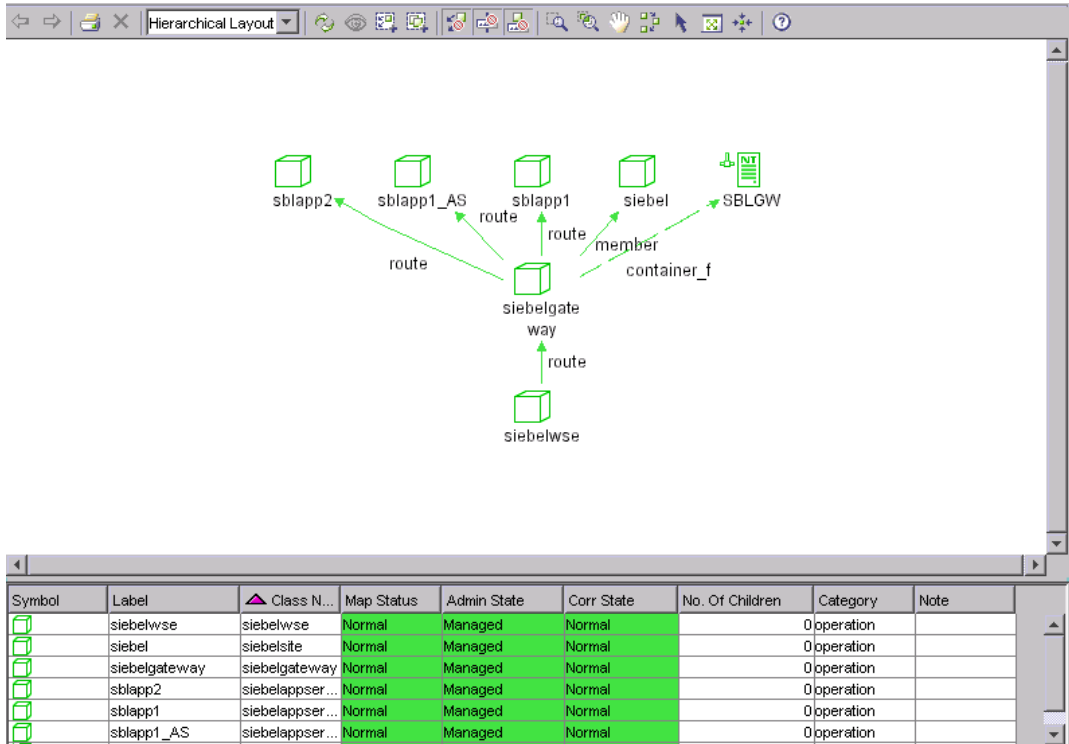
3 (Optional) To add to the displayed map interdependency information about one of the connected objects, right-click the object whose interdependencies you want to display. Select **Get Neighbors** and one of the following commands:

➤ **Next on Layer –** (Enabled only after using the **Layer** option) Keeps the previous Layer display on the Map pane, and adds to that layer the connected objects of one of the displayed objects.

The following example shows how the Next on Layer command is applied to the SBLGW object, and consequently, its connected objects are added to the map:
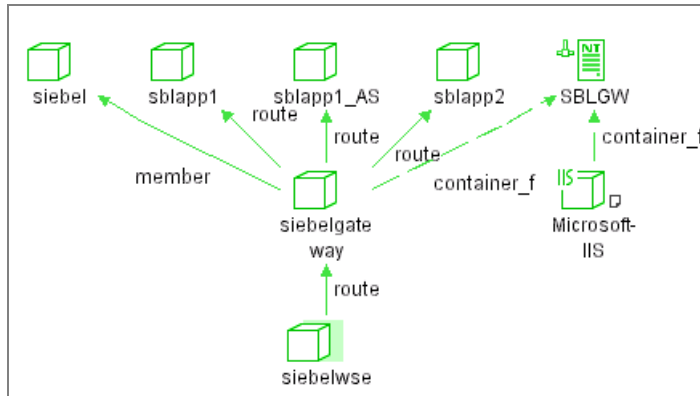
➤ **Next in View –** (Enabled only after using the **View** option) Keeps the previous View display on the Map pane and adds the connected objects from the given view of the selected object.
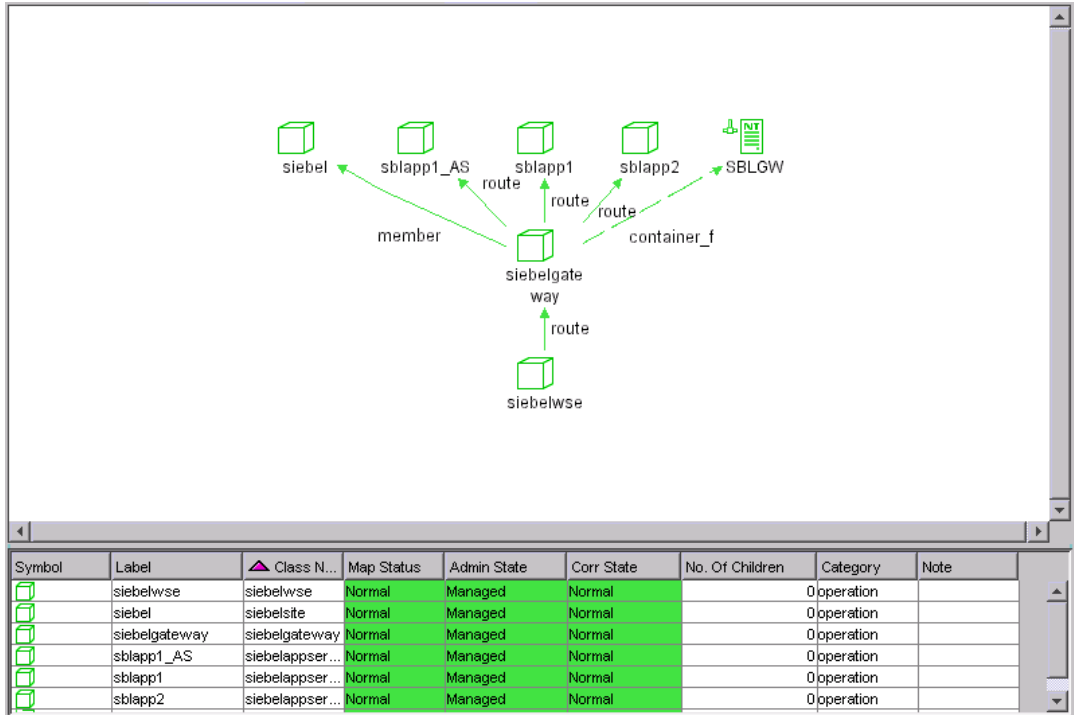
➤ **Next in Database –** (Enabled only after using the **Database** option) Keeps on the Map pane the previous **Database** display and adds to it the connected objects from the database of the selected object.
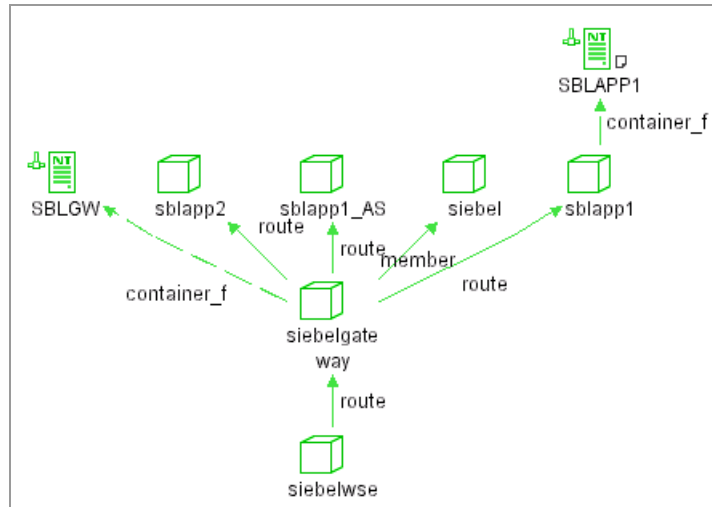
Showing a Correlation Impact

The **Show Impact** option enables you to select an object, which is defined by a correlation rule as a root cause object, and to display all the objects 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 objects that are affected by the IP object.

Note: For details on defining correlation rules, see Part VI, "Object Relationships."

To show the correlation impact of a root cause object:

1 In the Map pane or Information pane of Map View, right-click a root cause object and select **Correlation** > **Show Impact**.

> ➤ If one correlation rule is defined for the selected object, the Impact window is displayed. Skip to step 2.

> ➤ If several correlation rules are defined for this object, the Impact Correlation Rules dialog box is displayed.

> The Impact Correlation Rules dialog box lists all the correlation rules that define this object as a root cause object.

> Select a correlation rule and click **OK**.

2 The Impact window is displayed:
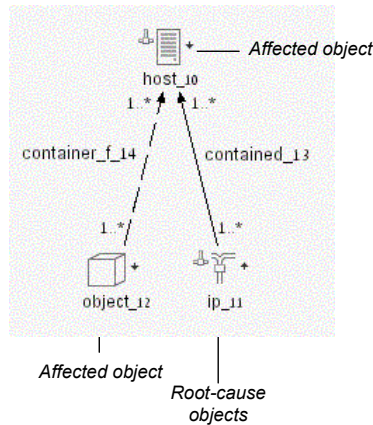
In the above example:

> ➤ The CPU object is the root cause object, and for that reason it is marked by this indicator.

> ➤ The HAVA object is part of the correlation rule, but it is neither the root cause object nor an affected one. For that reason it is marked by this indicator.

> ➤ The IP object is the affected object, and for that reason it is marked by this indicator.

Showing Root Cause Objects

Mercury Application Mapping enables you to select an object that is affected by a correlation rule, and display the root cause object(s) that affected it.

This section includes the following topics:

➤ "Showing the Root Cause Object of an Affected Object" on page 191

➤ "Problem Isolation" on page 192

Showing the Root Cause Object of an Affected Object

You can display the root cause object of an affected object.

To show the root cause object of an affected object:

1 In the Map pane or Information pane, right-click an affected object and select **Correlation** > **Show Root Cause**.

➤ If one correlation rule is defined for the selected object, the Correlation window is displayed. Skip to step 2.

➤ If several correlation rules are defined for this object, the Root Cause Correlation Rules dialog box is displayed.

The Root Cause Correlation Rules dialog box lists all the correlation rules that define this object as a root cause object.

Select a correlation rule and click **OK**.

2 The Correlation window is displayed:

191
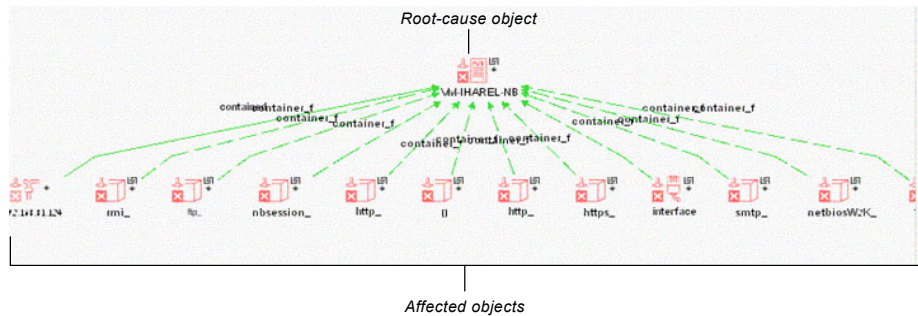
In the above example, the scenario that led to a correlation event occurring for the Switch object is displayed:

➤ The IP object on the left is the root cause object, and for that reason it is marked by this indicator.

➤ The objects at the bottom are part of the correlation rule, but they are neither the root cause objects nor the affected ones. For that reason they are marked by this indicator.

➤ The Switch object at the top is the affected object, and for that reason it is marked by this indicator .

Problem Isolation

This option enables you to retrieve root cause information for objects 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 classes that affect the object.

Show Problem Isolation, as opposed to **Show Root Cause**, allows you to view the source root cause object in one step. When using the **Show Root Cause** option to reach the root cause object, you have to use **Show Root Cause** on each trigger/affected rule, one by one.

To show the root cause object:

In the Map or Information pane, right-click an affected object and select
Correlation > Show Problem Isolation. A window, as illustrated below, is
displayed containing the route cause object and all the other classes in the
correlation rule chain.
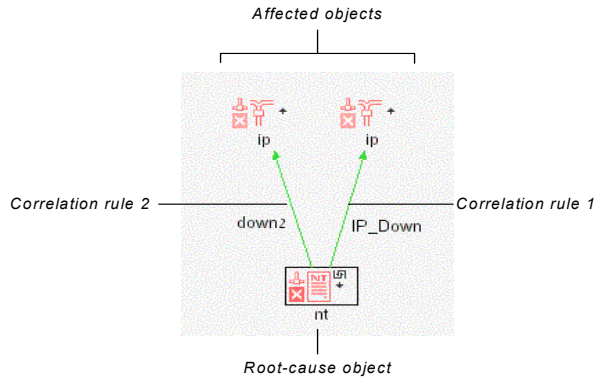
Use the **Show Problem Isolation** option on the IP object to view, in one step,
that the source of the failure is the Network object.

Use the **Show Root Cause** option to see that the host object is the failure as
it is the next class in the correlation rule chain. To display the Network
object as the root cause object, you have to right-click the host object and
select **Show Root Cause** again.

To view the information about specific objects that are triggered or affected
by the correlation rule, double-click the link connecting the objects. A map
containing the same information as when you use the Show Root Cause
option is displayed (for details, see "Showing Root Cause Objects" on
page 191).

Using Meta Links

A Meta Link is a logical link that represents several physical links between objects that are displayed on different layers. These links are not displayed together on one layer, due to the hierarchal structure of the view that organizes objects and their interdependencies in separate layers. The Meta Link option enables you to display in one window, the full path of physical links that sustain links at upper levels.

For example, in a certain view, a Client-Server link connects two NTs. However, the actual link 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 link:

➤ On the Map pane, double-click the Meta link of the contents you want to display. A new Link Map window appears, displaying the links that exists between the two Meta link's objects.

For example the following Meta link displays a logical route link between a host and a server:

SAINT

sql ——— *Double-click a Meta Link to open it*

TP-ALERTS
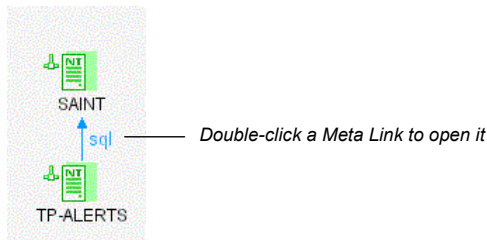
Double-click the route link to display all the objects that compose the physical path:
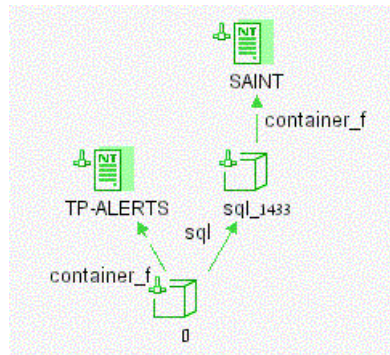
Note: If you double-click a link and a **Link Map** window is not displayed, it means that this link is not a Meta Link but a regular one.

Saving a Map View

You can save a Map View to a file.

To save a Map View to a file:

1 Select and display the Map View you want to save to file.

2 Select **Map** > **Export Graph** to open the Export Graph dialog box.
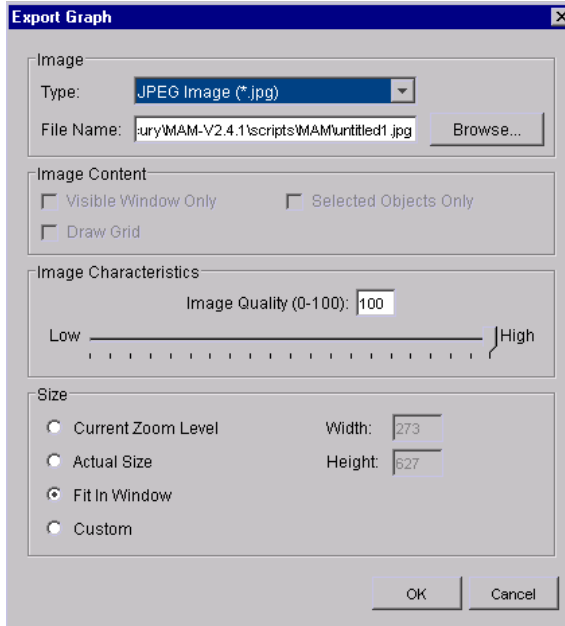
3 In the **Type** list, select the required file format.

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 Objects Only** to save only the objects that you selected and their links.

6 In the **Image Characteristics** area:

➤ Enter a value in the **Image Quality (0-100)** window or move the cursor on the scale. 100 means that the quality is excellent.

7 In the **Size** area:
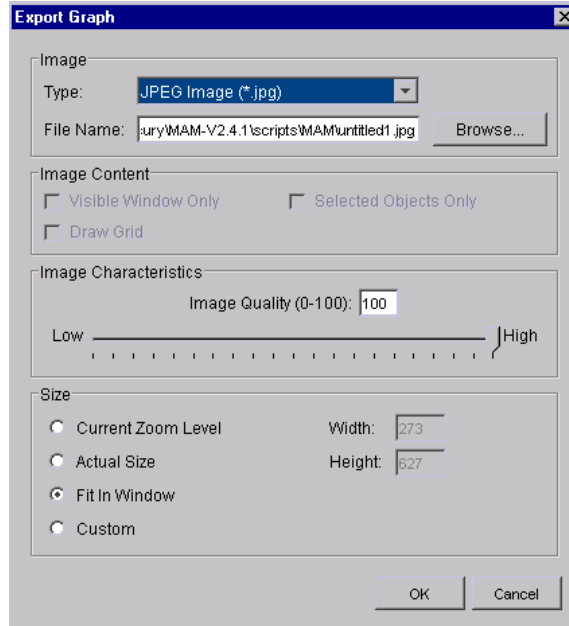
➤ Select **Actual Size** to save the view as it really is.

➤ 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**.

197

Printing the Map View

You can print the contents of the Map pane, 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 Map pane and Information pane tab according to your requirements, before printing.

This section includes the following topics:

➤ "Printing the Contents of the Map View in the Map Pane" on page 198

➤ "Printing the Contents of the Information Pane" on page 199

Printing the Contents of the Map View in the Map Pane

You can print the contents of the Map View in the Map Pane.

To print the contents of the Map View in the Map pane:

1 Select and display the contents you want to print in the Map pane.

2 To arrange the contents of the Map pane for print, you can use the **Layout**, **Layer Setup**, and **Hide** options.

For details about the **Layout** option, see "View Level" on page 41.

For details about the **Layer Setup** option, see "Improving Views" on page 172.

For details about the **Hide** option, see "View Shortcut Menu" on page 138.

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 links.

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 200

➤ "Hiding a Column" on page 200

➤ "Customizing a Display" on page 201

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 Object 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 Object 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 Objects

Mercury Application Mapping searching capabilities enables you to search for objects in the Map View, either in specific views and layers or in the entire CMDB. These capabilities also provide different search criteria, through which you can search for objects according to their class, label, and/or attributes.

To search for objects:

1 (Optional) If you want to search in a specific layer, open it in the Map View.

2 From the **Edit** menu, select **Find** to open the Find dialog box.

3 From the **Class** list, select the class of the object(s) you want to find.

4 (Optional) Select **Derived** to find and display both the selected class object(s) and any objects derived from the selected class (inheritance objects).

5 (Optional) In the **Options** area, if you want to find a specific object(s) whose label you know, enter the object'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 an object(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:

• **Attribute Name** – Select the required attribute from the list.

• **Operator** – Select the required operator (such as =) from the list (for details, see "Operator List" on page 45)

• **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 links.

➤ 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:

> ➤ **Current Layer** – Searches only in the layer that is currently displayed in the Map View.

> ➤ **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 CMDB.

8 (Optional) By default, **Get Full Path** is selected, and the view layers that appear above the object(s) are displayed. Clear **Get Full Path** if you want to display only the view name in which the object(s) is located.

9 Click **OK**. The search results are displayed in the Find Results tab in the Information pane.

11

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 objects and links from the Mercury Universal CMDB or from an existing view.

What is an Instance View?

Instance views are views that are built by manually attaching individual objects and links directly to the view. Instance views are not based on TQL queries. You can attach any object from an existing view or from the CMDB into an instance view. If you choose to attach an object from the CMDB, all existing links between the selected objects are added to the view as well.

Creating an Instance View

You create a new instance view by selecting objects 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 an object or choose from a list of related objects.

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

This section describes how to create a new instance view and then attach the required objects to the view.

Note: You do not need to attach any objects 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 of **Map View**, 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 **Properties** tab in the View Properties dialog box. Define additional view properties as described in "Creating a View with the View Manager" on page 94.

6 (Optional) Click the **Folding Rules** button to open the Link Rules dialog box. Define link rules to define multiple link rules. For details, see "Defining Multiple Link Rules" on page 110.

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 object(s) that you want to attach to the view. You can select the objects you want to attach to the instance view in the following ways:

➤ In the **Map View** tab, select the required object(s) at different levels of the view hierarchy. You can select multiple objects by holding down the CTRL button and clicking the required objects. To view and edit an object's attributes, see "Displaying an Object's Attributes" on page 209.

➤ In the **Search** tab, search for objects in existing views or in the entire CMDB. For details, see "Searching for an Object" on page 209.

➤ In the **Related Objects** tab, select an object that is related to the selected object in the **Map View** tab. For details, see "Displaying Related Objects" on page 211.

9 Click the **Attach** button to attach the selected objects to the view. The attached objects appear in the **Object** box.

10 Click **OK** to create your view.

Note: It can take a few moments until the new view appears in the View Explorer.

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 an Object's Attributes

This section describes how to display and edit an object's attributes.

To display and edit an object's attribute:

1 In the Instance View Definition dialog box, select the **Map View** tab.

2 Select the object whose attributes you want to view.

3 Click the **Show Object's attributes** button to open the Object Attributes dialog box. The Object's Attributes dialog box enables you to view and edit the attributes defined for the selected object. For details, see "Viewing and Editing Object Attributes" on page 176.

Searching for an Object

Mercury Application Mapping searching capabilities enables you to search for objects 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 objects according to their class, label, and/or attributes.

To search for objects:

1 In the Instance View Definition dialog box, select the **Search** tab.

2 From the **Class** list, select the class of the object(s) you want to find.

3 (Optional) Select **Derived** to find both the selected class object(s) and any objects derived from the selected class (inheritance objects).

4 (Optional) In the **Options** area, if you want to find a specific object(s) whose label you know, enter the object'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.

5 (Optional) In the **Advanced** area, to search for an object(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:

- **Name** – Select the required attribute from the list.

- **Operator** – Select the required operator (such as =) from the list (for details, see "Operator List" on page 45)

- **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** – Select **And** or **Or** to link multiple conditions**.**

➤ After you complete the attribute definition, you can change it by using the **Edit** and **Delete** buttons.

6 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.

7 Click **Search**. The search results are displayed in the **Search** tab.

8 To do another search, click **Back to Search**.

Displaying Related Objects

Mercury Application Mapping enables you to display the interdependencies of any selected object (for details, see "Displaying Interdependent Objects" on page 181). When you create a new instance view, you can choose to populate the view with objects from a list of related objects.

You can display:

➤ All objects connected to a selected object that are part of a specific view.

➤ All objects connected to the selected object in the entire CMDB.

You select an object in the **Map View** tab, and the objects related to the selected object appears in the **Related Objects** tab.

To display the related objects of a selected object:

1 In the Instance View Definition dialog box, select the **Map View** tab.

2 Select the object whose related objects you want to display.

3 To display related objects, do the following:

➤ Click the **Get Related Object from View** button to display related objects in the selected view.

➤ Click the **Get Related Object from Database** button to display related objects in the CMDB.

The related objects are displayed in the **Related Objects** 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 207.

12

Event Management

This chapter introduces the Event system, which monitors the changes occurring in the managed world and delivers information about them.

About 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 objects. Three of the Mercury Application Mapping sub-systems—Discovery, Correlation, and Time Rules—determine which kind of changes that occur in object attributes are registered as raw events.
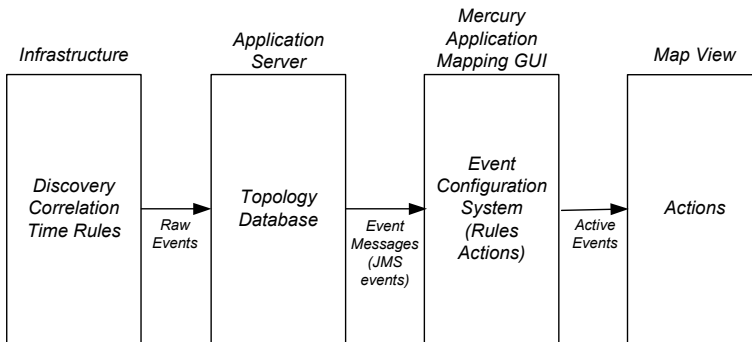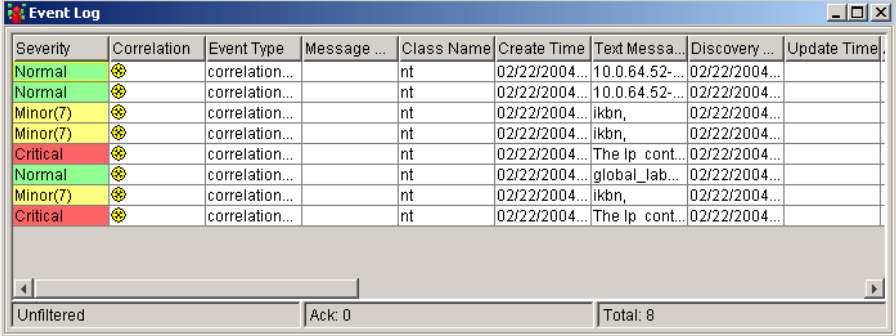
This section includes the following topics:

➤ "Raw Events" on page 215

➤ "Viewing Event Summary in the Status Bar" on page 218

Raw Events

Raw events are saved to the CMDB. If a raw event is related to an object that is displayed in the Map View, the raw event appears in the object's event log. For further details on the Event Log, see "Event Log" on page 228.

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 Ip 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 Ip 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 objects 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 the *Mercury Application Mapping Administration Guide*.

215

Example

The discovery system has determined that the administrative state of a certain object 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 an object, it should trigger an action that changes the object's admin state in the Map View.

Since the new raw event meets this rule's condition, the rule's action is performed, and the admin state of the object is changed accordingly..
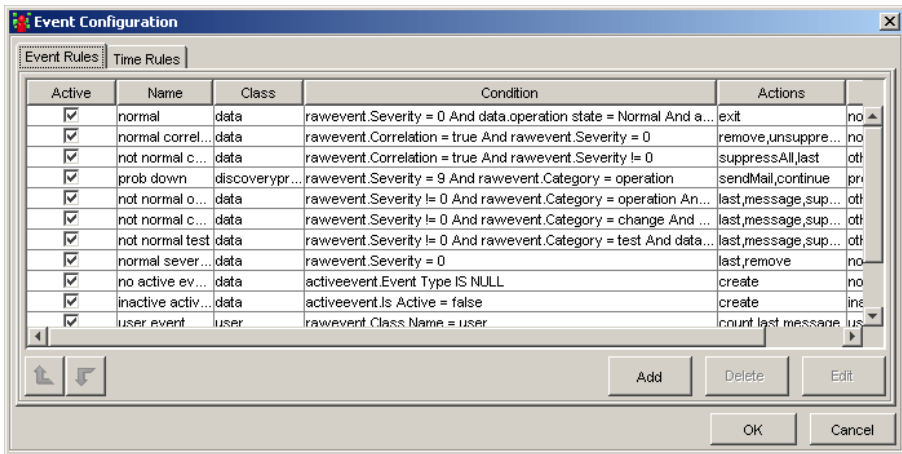
Although the Discovery, Correlation, and Time Rules systems continuously create raw events in response to changes occurring in the managed world, only those that meet at least one of the rule conditions are actually reported in Mercury Application Mapping. However, not all the raw events that meet the defined conditions are displayed in the Event browsers and manifested via state changes and icon blinks in the Map View. 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 Map View to manage active events related to the objects displayed in a selected view. This gives you the ability to locate problems that may be occurring in the system at the object level. When an object contains an event, except for an event that is a result of a correlation rule, the icon for the selected object blinks and an exclamation mark is added before the icon in the Explorer pane. You can then select the Object Events tab, or the other Event tabs and browsers, to view the event information and acknowledge the event:

Acknowledged event

Object Browser	Object Events	All Events	Inbox	View Events					
Severity	Correlation	Event Type	Class Name	Text Messa...	Discovery ...	System Ti...	Acknowled...
Critical		general_sy...	ip	State was c...	02/22/2004...	02/22/2004...	☑
Minor(4)		general_sy...	nt	State was c...	02/24/2004...	02/24/2004...	☐

Active events in the Object Events tab

Unacknowledged event

Once an event is acknowledged, the related icon on the Map pane stops blinking and the exclamation point disappears from the Explorer pane.

Viewing Event Summary in the Status Bar

You can view the summary of an event in the status bar.

| Unfiltered | Ack: 0 | Total: 1 |

The Status bar at the bottom of the Information pane, provides a quick way of viewing the following:

➤ Whether events are being filtered.

➤ The number of events that have been acknowledged.

➤ The total number of current events for the selected view.

The Event Tabs in the Information Pane

The Information pane contains the following tabs, enabling you to manage events either at the object or view level. For details of the attributes, see "Event Tab Definitions" on page 220.

Note: Event logs and browsers (regular and compound) display not more than 500 events.

This section includes the following topics:

➤ "Object Events" on page 219

➤ "View Events" on page 219

➤ "All Events" on page 219

➤ "Inbox" on page 219

➤ "Event Tab Definitions" on page 220

Object Events

The Object Events tab displays all the events related to the object currently selected in the Explorer or Map pane. It includes information about the attributes of each listed event.

View Events

The View Events tab displays all events related to objects of a selected view in the Explorer pane. It includes information about the attributes of each listed event.

All Events

Tip: Double-clicking an event in the All Event tab takes you to the related object in the Explorer pane and in the Map View.

The All Events tab displays all the events that occurred in relation to the system's views. It includes information about the attributes of each listed event.

On the All Event tab, you can identify to which view(s) an event is related, by double-clicking it and opening the Event Navigator dialog box.

The Event Navigator displays a list of the related views, and allows easy navigation between them. By selecting the view that is of interest to you, you can open it and examine the context of the event.

Inbox

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.

Event Tab Definitions

The following table describes the attributes of Mercury Application Mapping active events that make up the Event 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.
Event Type	The type of the event.
Message Group	Not in use.
Class Name	The class name of the event's object.
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 object 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.
Attribute Value	For events from Collector Threshold type only. Indicates the attribute value that caused the event.

Column	Description
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.
Object Global Label	The inclusive label of the object related to the event. Includes other objects that identify the selected object. 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 "Object Shortcut Menu" on page 140), 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.
Counter	The number of times the same active event has arrived without any severity or message change.

Column	Description
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 object'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 object, the icon still blinks.) The suppression action automatically defines, for all correlation rules, that objects that are affected by a root-cause object do not blink even though they have active events. For example, if a correlation rule is defined for the connection between a network object 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 object, 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, as described in "Defining Nodes, Links and Groups" on page 98).
	For example, if the severity of the event is **5** and the status weight is defined at **100%**, the object's Event Status is **5** (5*100/100). However, if the weight is defined as **60%**, the object's status is **3** (5*60/100).

Column	Description
Category	The category for which the event is sent. For details on categories, see the *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 an object.
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 object.
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 object(s) that affected the selected link. If only one correlation rule is defined for this object, the Correlation window is directly displayed.
Show Object Attributes	Displays the attributes of the object related to the selected event. You can edit the values of the attributes that are displayed in italics.
Event Navigator	Display the Event Navigator dialog box, which lists the node(s) to which the event is related. By double-clicking a node, you can open it and examine the context of the event.

Option	Description
Filter	(This option is enabled only if Event Filters are defined).
	Select one of the following options:
	➤ **Filter Names –** (Enabled only if you already used filters in your current session.) Displays the names of the last used filters. By clicking a name, you apply the selected filter to the event tables.
	➤ **Clear –** (Enabled only when a filter is applied on the event tables.) Enables you to remove the applied filter, and to display the event tables in their entirety.
	➤ **More Filters –** Displays a list of the available filters. By selecting a filter from the list, you apply it to the event tables.
Charts	Select one of the following options:
	➤ **Filters by Row**
	➤ **Chart Names** – Displays the names of the last used charts. By clicking a chart name, you display the selected chart, which refers to the data of the table from which you opened it.
	➤ **Chart Manager –** Displays the Chart Editor dialog box, enabling you to display one of the existing charts or to create new ones. For details, see Chapter 13, "Generating Charts."

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 226

➤ "Hiding a Column in an Event Tab" on page 226

➤ "Displaying Hidden Columns in an Event Tab" on page 226

➤ "Customizing an Event Tab" on page 227

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 227)

➤ Event Log (for details, see "Event Log" on page 228)

➤ Compound Event Browser (for details, see "Compound Event Browser" on page 228)

➤ Compound Event Log (for details, see "Compound Event Log" on page 229)

Event Browser

The Event browser displays the same information contained in the Event Browser tab in a separate window.

To display the Event Browser in a separate window:

Right-click an object in the Explorer or Map pane and select **Events** > **Event Browser** to open the Event Browser window.

For details about the information displayed, see "Event Tab Definitions" on page 220.

Event Log

The Event log displays a list of raw events that occurred to a selected object. Through this log, you can view the event history of the object.

To display the Event log:

Right-click an object in the Explorer or Map pane 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 220.

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 object's children (if children exist).

To display the Compound Event browser:

Right-click an object in the Explorer or Map pane and select **Events** > **Compound Event Browser** to open the Compound Event browser.

For details about the information displayed, see "Event Tab Definitions" on page 220.

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 object's children (if children exist).

To display the Compound Event log:

Right-click an object in the Explorer or Map pane and select **Events > Compound Event Log** to open the Compound Log browser.

Setting and Using Event Filters

Filters can be created to reduce the number of events displayed in the Event tabs, browsers, and logs. Those events that meet the criteria defined by the selected filter are displayed, enabling you to focus on events of particular interest and significance. You can define multiple filters, but only one filter can be active at a time.

Note: Event filters can be defined if **Show Events** is selected in the Create New View dialog box. For details, see "Creating a View" on page 92.

The Mercury Application Mapping Filter Editor automatically saves all created filters in the database—Browser, Tab and Log filters—thus enabling the sharing of filters among users. Users can select the filters that they want to use in the Events Filters and Event Log Filters tabs in the User Profile dialog box (for details, see "Understanding Layer Layouts" on page 154).

Defining an Event Filter

You can create filters to reduce the number of events displayed in the Event tabs, browsers and logs.

To define an event filter:

1 From the **Edit** menu, select **Filter Editor** to display the Filter Editor.

➤ To define a filter for the Event tabs and browsers, select the **Event** tab.

➤ To define a filter for the Event Log, select the **Log** tab.

2 Click **Add** to create a new filter.

3 In the **Name** box, enter a unique name for the filter.

4 (Optional) In the **Description** box, enter a description of the filter.

5 Click the **Add** button to open the Filter Restriction dialog box.
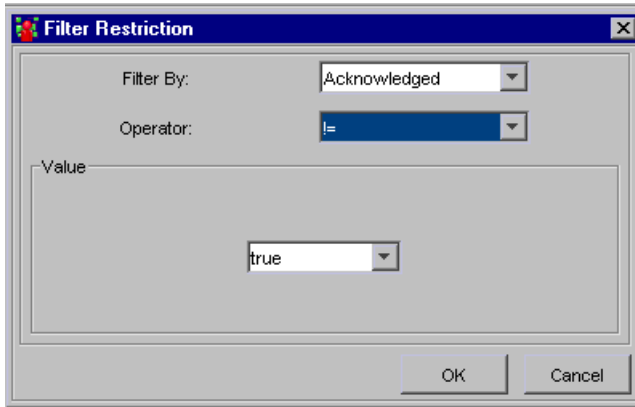
6 The **Filter By** list contains the event's attributes (for details, see "Event Tab Definitions" on page 220). From the list, select the attribute you want to use as the filter restriction.

7 From the **Operator** list, select an operator for the attribute you selected (for details, see "Operator List" on page 45)

8 In the **Value** box, select or enter a value, to which the operator should apply.

For example, to display only unacknowledged events, select the **Acknowledged** attribute. Define the operator as **not equal to** and the value as **true**.

Click **OK** to add the filter restriction to the filter. You return to Filter Editor dialog box.

9 (Optional) To add restrictions to the filter, repeat steps 5 to 8. If there are several restrictions, click in the **AND/OR** column to define their relationships as follows:

Select **And** to display those events that match both filter conditions or select **Or** to display those events that match either of the filter conditions.

10 To save the filter you defined, click **OK** in the Filter Editor dialog box.

To edit an existing condition:

Select the required condition in the Filter dialog box and click **Edit**.

To remove an existing condition:

Select it in the Filter dialog box and click **Remove.**

To delete an event filter:

Right-click an event related to the filter in the All Event or Object Event tabs in the Information pane and select **Clear Filter**.

Using an Event Filter

You can use existing filters to reduce the number of events displayed in the Event tabs, browsers and logs.

To use an event filter:

1 Open the Event browser, tab, or log you want to filter.

2 Right click the object and select **Filters** > **More Filters** to open the Filters dialog box.

Note: If you have used filters in the current session, the list of used filters also appears on the shortcut menu. You can select a filter from the menu.

---

3 Select the filter you want to use

4 Click **OK**. The selected filter filters the Event table.

To remove an event filter:

1 Open the Event browser, tab, or log from which you want to remove the filter.

2 Right-click the object and select **Filters** > **Clear**.

Receiving Change Events

This section describes the Change event type. It includes:

➤ "Receiving Change Events When Objects Are Added or Deleted" on page 232

➤ "Receiving Change Events When a Change Monitored Attribute is Modified" on page 233

Receiving Change Events When Objects Are Added or Deleted

Mercury Application Mapping enables you to receive events of the type Change when:

➤ objects are added to maps

➤ objects are deleted from maps

➤ Static attributes are modified

To receive events of type Change when objects are added to, or deleted from, maps:

1 In the View Manager, right-click the required object and then click **Properties** to open the Properties tab in the View Properties dialog box.

2 In the **Send notification for the following changes:** section select:

➤ **Added Objects**

➤ **Removed Objects**

3 Click **OK** to save the changes you have made. An event is sent every time an object 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 class 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 Class Browser, ensure that **Change Monitored** is selected for the attribute in the Class Attributes page.

2 In the View Manager, open the View Properties dialog box.

3 Click the **Advanced** tab and add the categories.

For details on using the Advanced tab, see step 2 in "Defining Map Categories" on page 95.

To view change events and statuses on the maps themselves, add the Change category.

4 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 object has changed, the name of the object that has been changed is displayed.

Note: Even if the object's color does not change due to the priority you assigned it in the **Category Priority** list, the event still appears in the All Events tab.

Sending an Event

This section describes how to define the event category that affects an object.

To send an event:

1 Right-click an object in the Map 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 Enumeration Manager. For details, see Chapter 26, "Enumerations and Lists."

3 Select the required severity.

4 Click **OK**. The color and blinking of the object are changed when the event is sent under the conditions defined in "Defining Map Categories" on page 95.

Clearing All Events

This option enables you to delete the events belonging to an object of a certain category or of all categories.

To clear the events from an object:

1 In the Map View, right-click the object 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 object may still be affected by other category events. Only if you choose to delete the events for all category types, does the object's state return to Normal.

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

Part V

Generating View Indicators and Charts

13

Generating Charts

This chapter describes how to generate charts.

About Generating Charts

Charts are an additional visual representation of views and objects. Like System reports (for details, see "Creating System Report Templates" on page 335), charts enable you to produce statistics about the network's state, and to view comparisons, tendencies, and trends in the managed data.

The data that is displayed in Mercury Application Mapping tables can be represented in charts. The tables that can be converted into charts include the tables that appear in the Information pane, Event logs and Browsers, and System Reports.

Charts are based on data series of one or more table's columns. You use the Mercury Application Mapping tools to determine the scope of the data series that are displayed in a chart. To narrow the selected data series, you can hide a table's rows, thus focusing on certain rows that contain important information. On the other hand, to add data to a chart, you can use the Chart Manager to create layered charts, in which each layer contains additional information about the layer above it.

Charts are automatically saved in the Chart Manager. You can share charts with other users, you can repeatedly use the same chart to detect changes in the managed data, and you can reuse one chart for the representation of different views.

If you have more than one chart you can change their order in the report by selecting one of the charts and clicking:

➤ The left button to move the chart up.

➤ The right button to move the chart down.

Converting a Report into a Chart

You can use a report to create a chart that graphically display the view's data. This capability is available only if you create a report in table format.

To convert a report into a chart

1 Create a report in a Table format and display it – for details, see "System Reports" on page 327as described in "Introduction to the Report Manager" on page 329 or "Configuration Management Reports" on page 263.

2 Right-click the table and select **Chart Manager** to open the Chart Editor dialog box.

3 Click the **Add** button to create a new chart based on the report data, and follow the steps described in "Defining a Chart" on page 241.

Defining a Chart

To add charts to a report you must first define a chart and then select the type of chart you want to create: pie, column, or line.

For example:
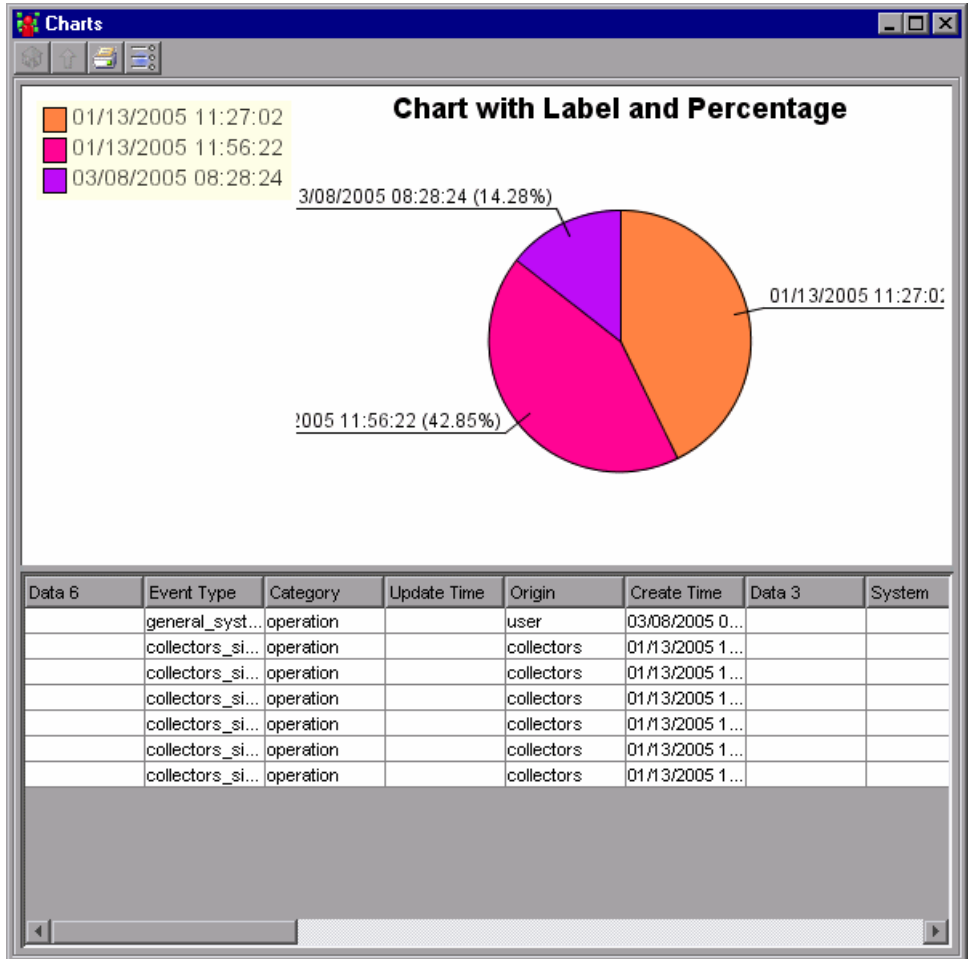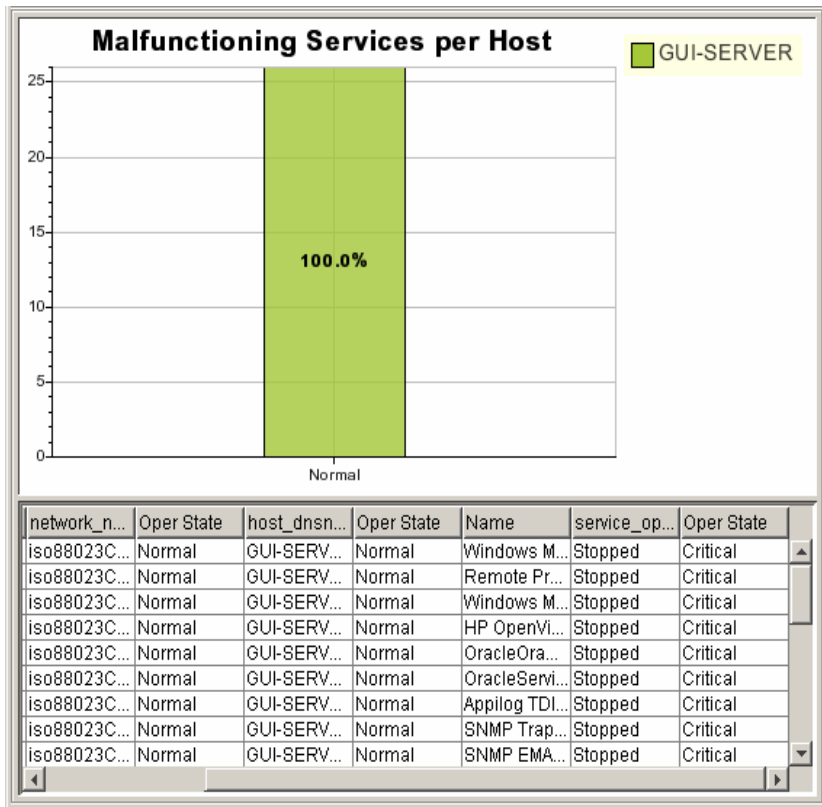
To define a chart:

1 Right-click the table of the data you want to display as a chart and select **Chart Manager** to open the Chart Manager dialog box.
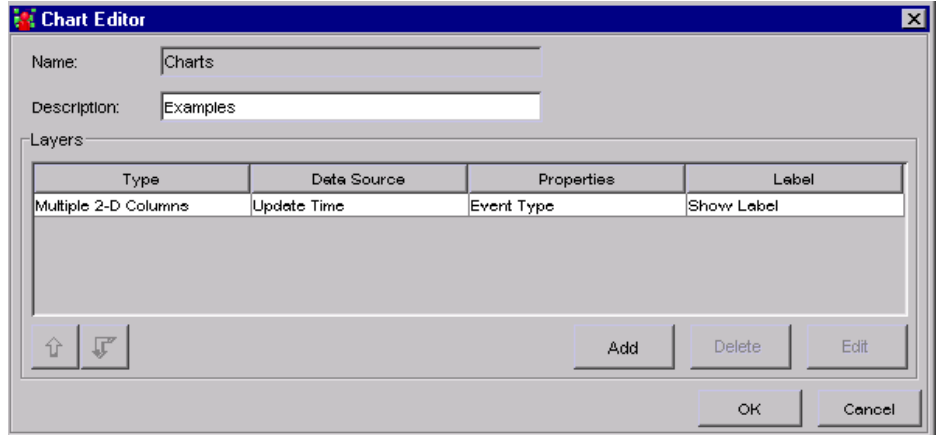
2 Click the **Add** button to open the Chart Wizard.
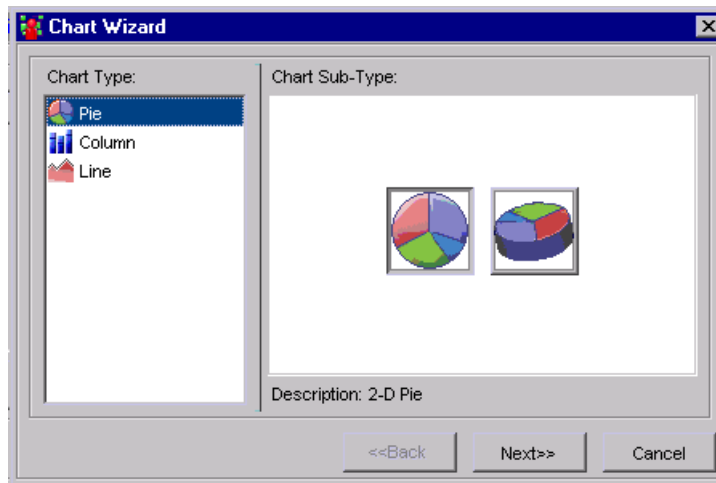
3 Click **Chart type** area, select a chart type:

➤ **Pie** (for details, see "Creating a Pie Chart" on page 244)

➤ **Column** (for details, see "Creating a Column Chart" on page 245)

➤ **Line** (for details, see "Creating a Line Chart" on page 246)

Creating a Pie Chart

Pie charts display the contribution of each value to a total and compare various components of one entity.

You can choose from the following pie charts:

➤ **2-D Pie** – a 2-dimensional pie chart

➤ **3-D Pie** – a pie chart with a 3-dimensional visual effect

Mercury Application Mapping pie charts can be marked with object labels, values, and/or percentages.

To create a pie chart:

1 Access the Chart Wizard, select **Pie** in the **Chart type** list, select a sub-type in the **Chart Sub-Type** box, and click **Next**.

2 Select the table column whose data series you want to present in the **Data Source** list. This list contains the column headers of the table you selected to display as a chart.

3 Click **Segmentation** to count the number of instances that are of the same type in the selected column and displays their division.

For example, if you open the Chart Manager dialog box from the Object Browser tab, and select **Oper State** as the data source, selecting Segmentation will graphically display the operational state segmentation of the managed objects.

4 Click **Value** (Integer type data only) to compare and display the numeric values of components of one entity

For example: To display the frequency of events per object, open the Chart Manager dialog box from the All Events tab, select the **Counter** as the Data Source, and select **Label** from the **Label** list.

5 Click **Scale (Divide by)** (Boolean, Date, or Number—Integer or Double—type only) to divide the data you select into several groups, according to the number of groups you specify (between 1 and 9). If you select data from Date type as the source data, the **Scale** box toggles to date buttons.

For example, to display the status division of the system events, open the Chart Manager dialog box from the All Event tab, select the **Event Status** as the Data Source, and specify a number in the **Scale** box.

6 Click **Next**.

7 In the **Chart Title** box, enter a title for the chart. This title appears at the head of the chart.

8 From the list of option buttons, some of the following types of data labels can be displayed depending on the type of chart:

> ➤ Select **None** to display the chart without labels.

> ➤ Select **Show Label** to display the chart with labels.

> ➤ Select **Show Value** to display the chart without labels but with values.

> ➤ Select **Show Percentage** to display the chart without labels but with percentages.

> ➤ Select **Show Relative Percentage** to display the chart without labels but with relative percentages.

> ➤ Select **Show Absolute Percentage** to display the chart without labels but with absolute percentages.

> ➤ Select **Show Label and Percentage** to display the chart with labels and the values in percentages.

> ➤ Select **Show Label and Value** to display the chart with labels and with values.

9 Click **Finish**.

Creating a Column Chart

Column charts compare values across categories and are used to describe trends in categories of data.

> ➤ **2-D Column** – clustered column which displays one data series or compare multiple series

> ➤ **3-D Column** – column with 3-dimensional visual effect

> ➤ **Stacked Column** – compares the contribution of each value to a total across categories. Available for 2-D and 3-D columns.

Charts with one column's data can be marked with object labels, values, and/or percentages. Charts with stacked or multi column data can also display the absolute or relative percentage of the values.

To create a Column chart:

1 Access the Chart Wizard, select the **Column** chart type, select a sub-type, and click **Next**.

One of the following boxes is displayed:

➤ If you select the subtype that is based on one column's data, the dialog box that opens is the same as for pie charts. For details, see "Creating a Pie Chart" on page 244.

➤ If you select stacked or multi column subtype to use two columns of data, the Segmentation Count dialog box is displayed.

The **Data Source** and **Group by** boxes contain the column headers of the table you selected to display as a chart.

- **Data Source –** select the table column whose data you want to display

- **Group by –** select the table column according to which you want to sort the source data

For example, to display the segmentation of map status categories in each class, open the Chart Manager from the Object Browser tab, select **Map Status** as the Data Source, and **Class Name** in the **Group by** list.

2 Click **Next** to define the chart title and labels.

3 In the **Chart Title** box, enter a name for the chart. This title appears at the head of the chart.

4 From the list of option buttons, select the type of data labels for the chart.

5 Click **Finish**.

Creating a Line Chart

Line charts display trends over time or over categories, and are used to describe data trends that can be perceived at a glance.

You can choose from the following line graphs:

➤ **X,Y Line** – A line with markers displayed at each data value

➤ **Area Line** – A stacked line

➤ **Ribbon Line** – A line with a 3-D visual effect

➤ **Ribbon Area Line** – A stacked line with a 3-D visual effect

Line charts can use numbers or time as a measurement value. Area charts can be used for many of the same tasks as line charts, but are typically used to describe cumulative totals, with each area corresponding to one part of an entire trend.

To create a line chart:

1 Access the Chart Wizard, select the **Line** chart type, select a sub-type, and click **Next**.

2 In the **Axis** area:

➤ Select the table column whose values you want to display as the X-axis values in the **(X) axis values** list (Only for data of Integer or Date types).

➤ Select the table column whose values you want to display as the Y-axis values in the **(Y) axis values** list.

3 In the **Line Instances** area:

➤ Select the column whose instance(s) you want to display in the chart in the **Column name** list.

➤ Select one or more instances for display in the **Instance** box. This field contains all the existing instances of the column you selected.

4 Click **Next**.

5 Enter the name of the chart in the **Chart Title** box.

6 Enter the X-axis title in the **Title (X) axis** box.

7 Enter the Y-axis title in the **Title (Y) axis** box.

8 In the Data Labels area, select one of the following:

➤ Select **None** to display the chart without labels.

➤ Select **Show X Value** to display the chart with the values for the X-axis.

➤ Select **Show Y Value** to display the chart with the values for the Y-axis.

➤ Select **Show X, Y Values** to display the chart with the values for the X- and Y-axis.

➤ Click **Finish**. The chart is displayed (for details, see below).

Displaying a Chart

Once a chart is created, it is automatically saved in the Chart Manager dialog box for further use. Each time this chart is displayed, it refers to the currently selected view and reflects the most updated data. Thus, you can use the same chart to display data of different views and you can detect changes in data by re-displaying a chart from time to time in the course of your work.

To display a chart:

1 If the Chart Editor dialog box is not already open, select the view and table whose data you want the chart to reflect.

2 Right-click the table and select **Chart Manager**.

The list of charts that you already used in your current session also appear in the **Charts** submenu.

Note: If the chart you want to display appears on this list, you can open it directly by selecting it from the shortcut menu.

The Chart Manager dialog box is displayed.

3 From the list of charts, select the chart you want to display.

4 Click **OK** to display the chart.

Displaying Another Layer in a Chart

If you chart report has more than one chart, you can toggle between the charts (layers of information).

To display another layer in a chart:

1 With the chart open, double-click the segment whose underlying layer you want to display:
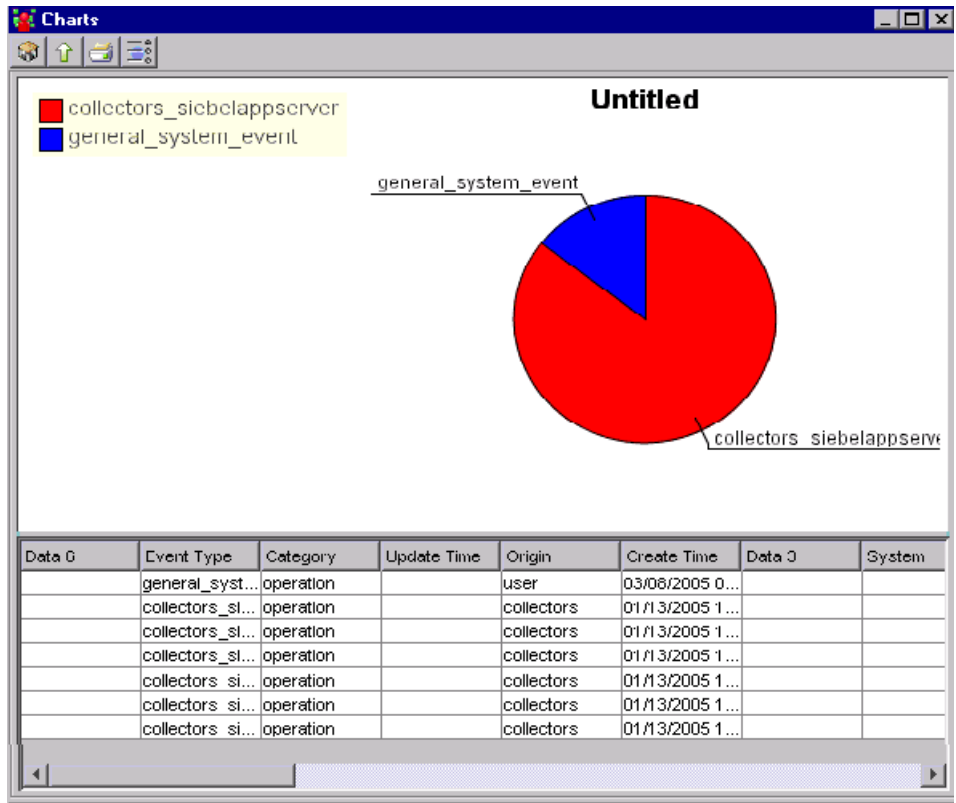
The corresponding segment of the lower layer is displayed.
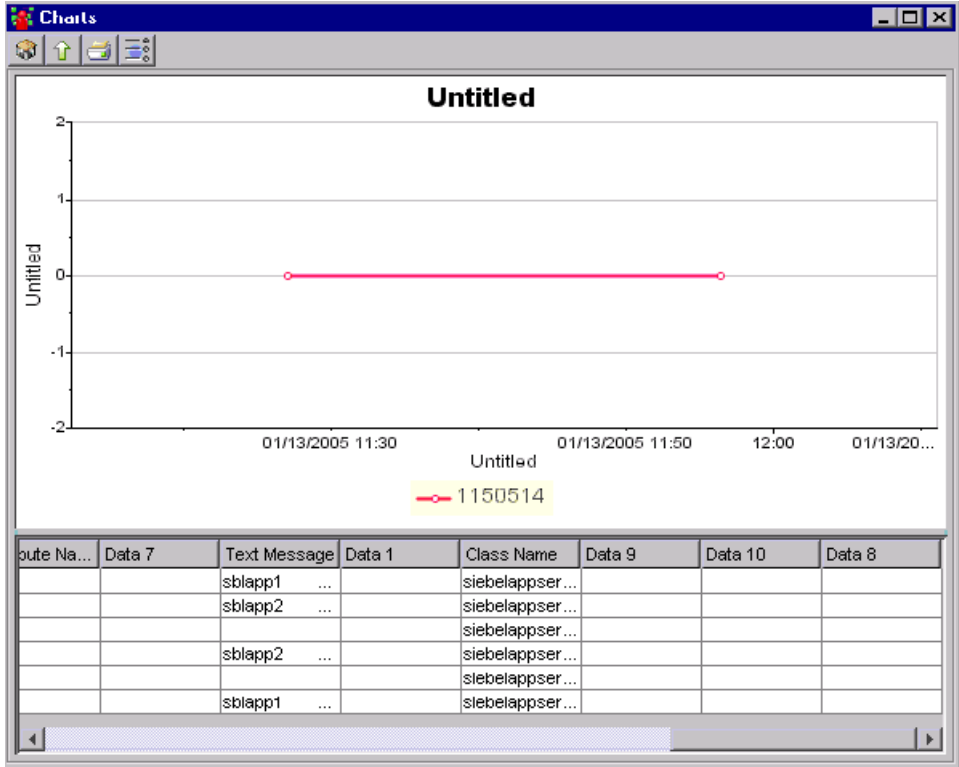
The lower level displays additional data about the selected segment in the higher level.

2 To move to the layer above the current one, click the **Up** button. This button appears only when you have displayed the lower level.

3 To return to the top level, click the **Home** button. This button appears only when you have displayed the lower level.

4 To edit the report properties, click the **Properties** button (for details, see "Editing an Existing Chart" on page 251).

5 To print the report, click the **Print** button (for details, see "Printing a Chart" on page 252).

Editing an Existing Chart

You can edit the properties of an existing chart, either provisionally for a one-time display, or permanently, for further use.

You can also permanently edit charts that you created.

To edit a chart for a one-time display:

1 With the chart open, click the **Properties** button on the chart toolbar or right-click anywhere on the Chart pane, and select **Properties** to open the Edit Chart dialog box.

2 Edit the available properties in the Data Source and Chart Label tabs. Once you have finished editing the chart properties, click **OK**. You return to the open chart, which reflects the changes you made.

To permanently edit a chart:

1 If the Chart Editor dialog box is not already open, display the table that relates to the chart you want to edit.

2 Right-click the table, and select **Charts** > **Chart Manager** to open the Chart Editor dialog box.

3 To edit the description of the chart, make changes in the **Description** box.

4 To change the order of the chart's layers, select the layer you want to move from the **Layers** list, and use the arrow buttons.

5 To edit other properties, select the layer you want to change from the **Layers** list, and click the **Edit** button to open the Edit Chart dialog box.

6 Edit the available properties in the Data Source and Chart Label tabs and click **OK**. You are returned to the Chart Editor dialog box.

7 To edit another layer, select it from the **Layers** list and repeat steps 5 and 6.

8 To add a further layer, click **Add** to open the Chart Wizard. Follow the instructions in one of the following:

> ➤ "Creating a Pie Chart" on page 244

> ➤ "Creating a Column Chart" on page 245

> ➤ "Creating a Line Chart" on page 246

9 To delete a layer, select the layer and click **Delete**.

10 To save your changes, click **OK**. You are returned to the Chart Editor dialog box.

11 To change the order of the layers after adding a layer, use the arrow buttons.

12 To display the chart you edited, select it in the Chart Manager dialog box and click **OK**. For details on accessing the Chart Manager, see "Displaying a Chart" on page 248).

Printing a Chart

You can print the chart diagram and/or the chart's table data.

To print a chart:

1 With the chart open, click the **Print** button. on the chart toolbar or right-click anywhere on the Chart pane, and select **Print** to open the Print dialog.

2 Select one of the following:

➤ To print the chart diagram, select **Chart** and click **OK**.

➤ To print the table that contains the chart data, select **Chart Data** and click **OK**.

14

Generating View Indicators

This chapter describes how to generate view indicators.

This chapter describes:	On page:
Understanding View Indicators	254
Changing Pie Indicator Size and Direction	261
Producing a View Indicator	262
Printing View Indicators	262

Understanding View Indicators

View Indicators enable you to quickly create and display dynamically changing diagrams that contain statistical data about several significant attributes of the system's views. View Indicators relate to data according to the way it is arranged and expressed in the Map View context. The View Indicators display the data graphically, dynamically reflecting the changes that occur in the data, and provide a fixed template for the data selection and presentation.

Each View Indicator relates to a specific attribute(s) of the view or the view's objects. When it is generated, it appears in a new and separate tab on the Map pane, and displays the results in a fixed diagram type (pie, bars, or dashboard). In each tab, selected data from all the existing views are displayed side by side, thus allowing you to compare the state and performance of your views:
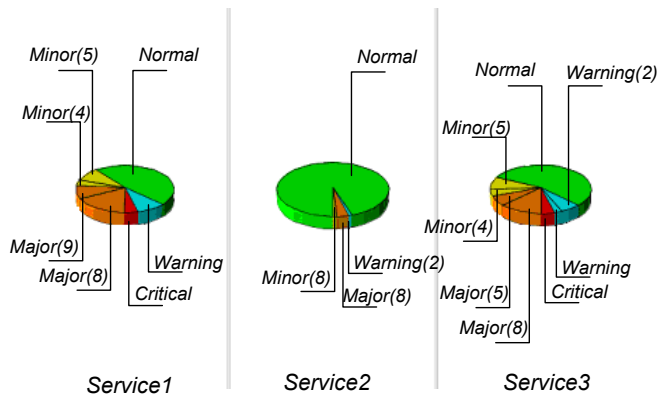
View indicators are based on data series, which appear in certain Mercury Application Mapping tables. In addition to graphically representing data, the View Indicators present statistics about the segments of the displayed data, in the form of pie and bar diagrams. These statistics specify the number of instances that are included in a segment, and the relative contribution of each segment to the total. You can see these statistics through a tooltip, by holding the cursor over a segment in the diagram:

Service 1
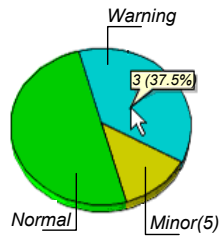
There are three View Indicators types: pies, bars, and dashboards. Each type is described in the following sections.

Note: The pie and bars indicators can be displayed in either 2-dimension or 3-dimension.

Pie Indicators

Pie indicators are based on one data series, that is, the data of one column. Choose between the following types (for details on creating pie charts, see "Creating a Pie Chart" on page 244):

➤ "Events by Status" on page 256

➤ "Events by Class" on page 256

➤ "Object Oper State" on page 257

➤ "Object Admin State" on page 257

➤ "Object Corr State" on page 258

Events by Status

A pie diagram that displays the segmentation of event's status in each view. This Indicator graphically reflects the Event Status column in the Event View tab on the Information pane.

Network
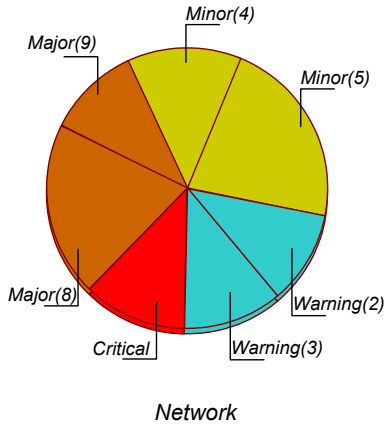
Events by Class

A pie diagram that displays the segmentation of event's classes in each view. This Indicator graphically reflects the Class Name column in the Event View tab:

Network
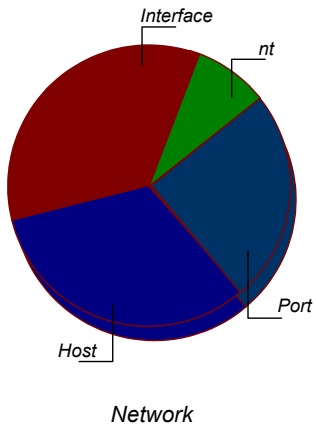
Object Oper State

A pie diagram that displays the segmentation of an object's operational state in each view. This indicator graphically reflects the Oper State column in the Object Browser tab:

Network
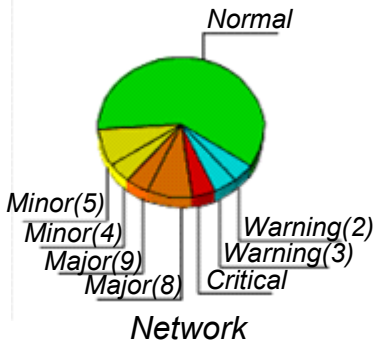
Object Admin State

A pie diagram that displays the segmentation of an object's Admin State in each view. This indicator graphically reflects the Admin State column in the Object Browser tab:

Network
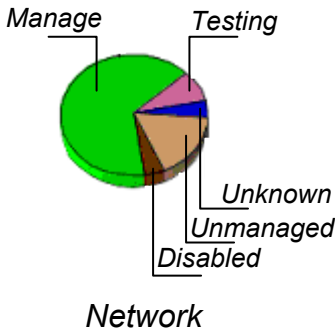
Object Corr State

A pie diagram that displays the segmentation of objects' Correlation State in each view. This Indicator graphically reflects the Corr State column in the Object Browser tab:
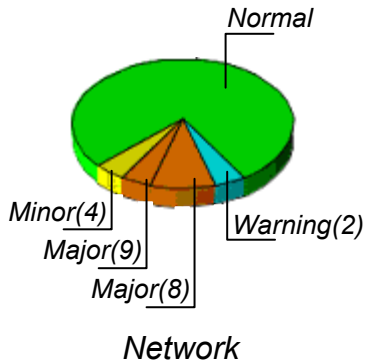
Bar Indicators

Bar Indicators are based on two data series, that is, the data of two columns. Choose between the following types:

➤ "Event By Class and Status" on page 258

➤ "Event Status and Ack" on page 259

Event By Class and Status

A Bar diagram that displays a comparison of event's status across classes. This Indicator graphically reflects the Event Status and Class Name columns in the Event View tab:
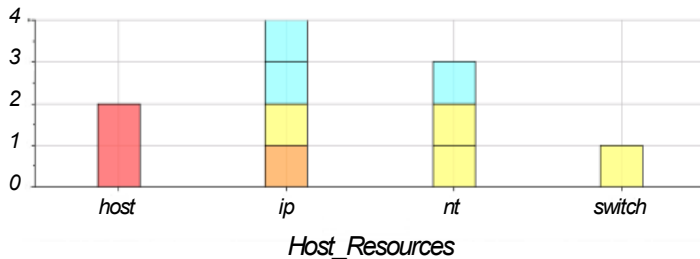
Event Status and Ack

A bar diagram that displays a comparison of event's status according to their acknowledgment state (true=acknowledged, false=unacknowledged). This indicator graphically reflects the **Event Status** and **Is Ack** columns in the Event View tab:
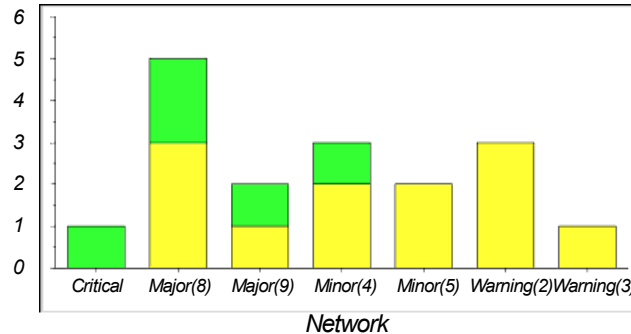
Dashboard Indicators

Dashboard indicators are based on availability calculations of views. These calculations are defined in the View Manager, and are set for each view separately. If no availability rules are defined for your views, the Availability Indicators is empty. You can see the data series that are represented in the Availability Indicators through System and Map reports.

The following dashboard indicators are available:

➤ "Daily Availability" on page 260
➤ "Weekly Availability" on page 261
➤ "Monthly Availability" on page 261

Daily Availability

A dashboard that displays view availability during one day, according to the settings you defined in the View Manager.

➤ The cursor in the dashboard points at a number that indicates, in percentages, for how long the view has been available, in comparison to the entire measured period. These statistics are also shown by the number at the center of the dashboard. In the following example, the view is available for **59%** of the day.

➤ The colors on the dashboard indicate the alert level that is assigned to the different unavailability durations. In the following example, when the view is available for 90% to 100% of the time, it is considered Normal (green); when it is available for 80% to 90% of the time, the alert level is Warning (light blue); when it is available for 70% to 80% of the time, the alert level is Minor (yellow); when it is available for 70% to 80% of the time, its alert level is Minor (yellow); when it is available for only 0% to 70% of the time, its alert level is Critical (Red).
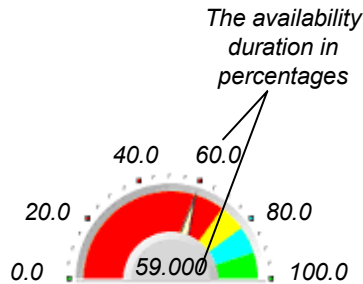
Network

Weekly Availability

A dashboard that displays the view's availability during one week.

Monthly Availability

A dashboard that displays the view's availability during one month.

Changing Pie Indicator Size and Direction

Once the pie indicators are generated, you can manipulate their size and direction as follows:

To change the size of a pie pane:

Use the **Zoom In** and **Zoom Out** buttons on the toolbar.

To increase the Pie's size:

Hold down **Shift** and drag the cursor from inside the pie outward.

To decrease pie size:

Hold down **Shift** and drag the pointer from outside the pie inward.

To rotate the pie:

Drag the pie in the required direction.

To change the pie location in its pane:

Right-click the pie and move it to the required location.

Producing a View Indicator

The indicators that are already displayed in the Map pane have a tick next to their name. You select the View indicator you want to generate. The View Indicator opens in a new tab on the Map pane.

To produce a View Indicator:

1 From the **View** menu, select **View Indicators** to display a list of the available view indicators (for details, see "Understanding View Indicators" on page 254).

2 The indicators that are already displayed in the Map pane have a tick next to their name. Select the View indicator you want to generate. The View Indicator opens in a new tab on the Map pane.

3 (Optional) You can print the displayed View Indicator, by selecting **Print Map** from the **Map** menu. For details, see the next section.

To close a View Indicator:

➤ Select **View** > **View Indicators** and click its name or right-click the View Indicator's tab on the Map pane to open a dropdown menu.

➤ To close only the selected indicator, click **Remove** <indicator name>.

➤ To close all the displayed view indicators, click **Remove All**.

Printing View Indicators

You can print the View Indicators data in the Map pane only. To print the availability data that is represented in the View Indicators, you must create Map Reports and print the data from them (for details, see "Generating a Map Report" on page 234).

To print a view indicator map in the Map pane:

1 With the View Indicator displayed on the Map pane, click **Print Map** on the **Map** menu to open the Print Setup dialog box.

2 Make the required changes and click **OK** to open the Print dialog box.

3 Click **OK**.

15

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 objects in a selected view and their attribute values.

To generate an Asset Report:

1 In the Map View, 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 **Class Attributes** tab in the Class Browser. For more details, see "Creating Classes" on page 374.

2 Select **Reports** > **Asset Report** to open the following in your browser.
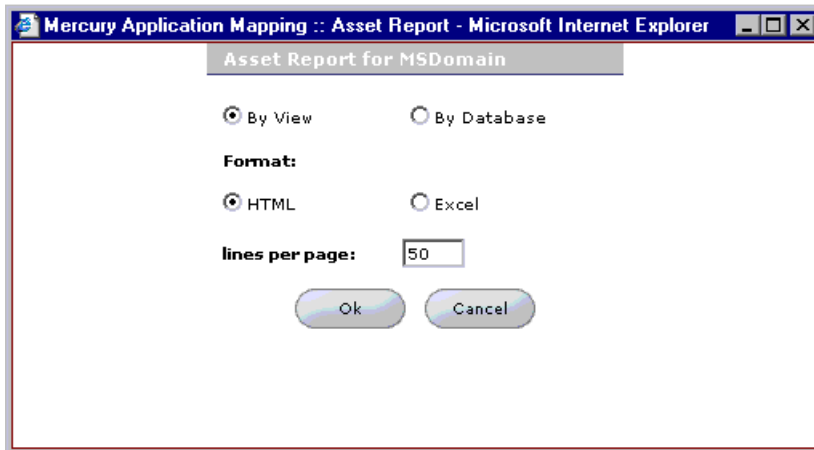
3 To define what you want the report to include:

➤ Select **By View** to include only objects contained in this view.

➤ Select **By Database** to include additional information from the Mercury Universal CMDB. By choosing this option, the report will include objects that are connected by the **container_f** link to the objects 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.
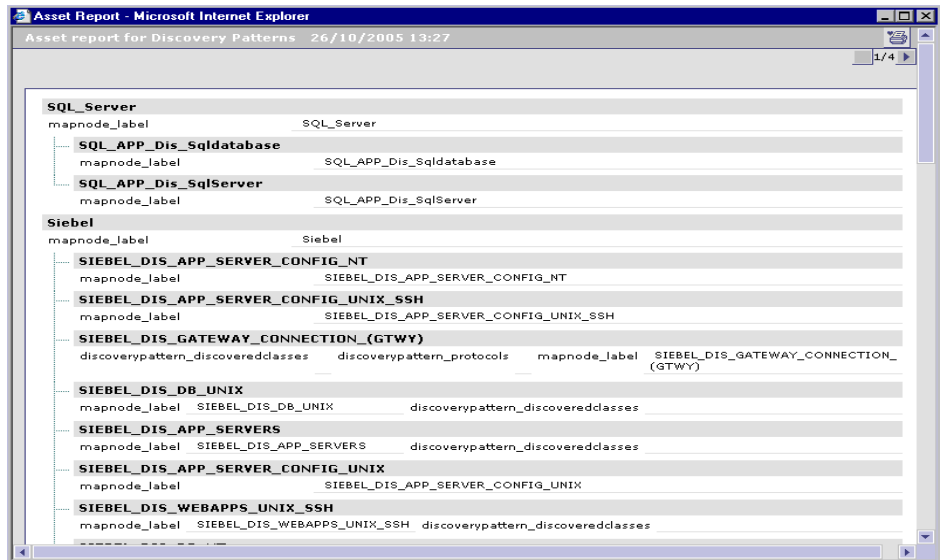
The Asset Report retains the same hierarchical structure as in the Map View.

Generating a Dependency Report

This report is based on TQLs that are created in the Logical Object Builder that use the link of the type **dependency**. Mercury Application Mapping enables you to create a report that lists all the links 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 links 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 Map View, right-click the view for which you want to create a Dependency Report.

2 Select **Reports** > **Host Dependency Report** to open the following in your browser.
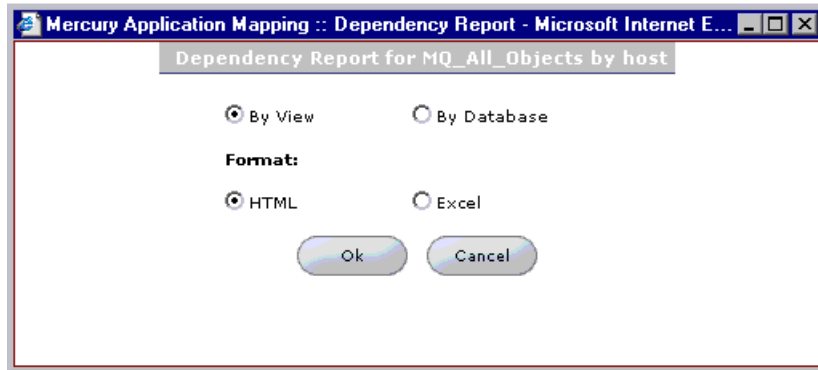
3 To define what you want the report to include:

➤ 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.

The report contains the following columns:

Column	Description
Host	The servers in the selected view.
Link	The type of link 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 Logical Object Builder using the **dependency** link to describe the connection between hosts. Every new **dependency** link 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 objects 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 objects whose events have been acknowledged.

➤ Select **False** to display objects whose events have not been acknowledged.

➤ Select **None** to display objects 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 objects whose events are the result of a correlation rule.

➤ Select **False** to display objects whose events are **not** the result of a correlation rule.

➤ Select **None** to display objects 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.
Class	The class name of the event's object.
Object	The object'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.

Creating and Displaying a Host Report

A Host report is a report that focuses on a certain host instance, and displays information and objects that are related to it according to the way they are stored and organized in the database.

Similar to a system report, a host report is based on a Report TQL. However, for the Report TQL to function as a host report, it must include a host object, and this host object must be the first in the node order of the report (as defined in the Report Manager), as shown in the following figure:
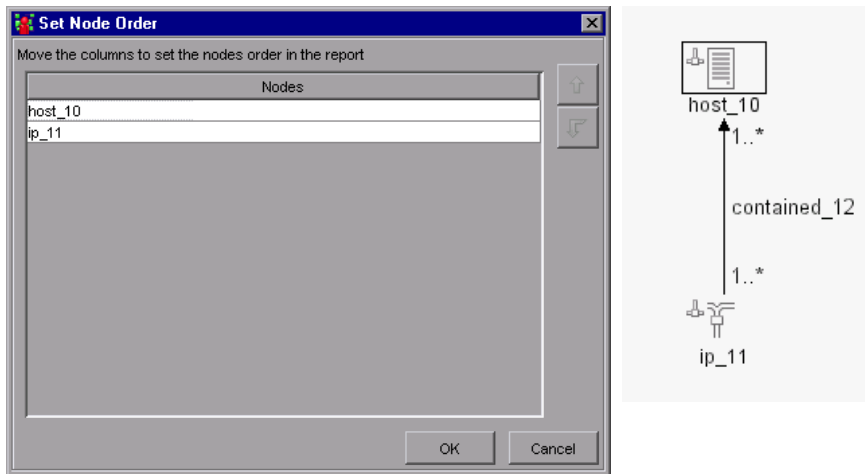
A host report is generated from the map view and displays the data of the selected object, in this case, a host object. Host reports relate to data as organized in the database and not in the current views.

To create a host report:

1 In the TQL Builder, create a Report TQL that includes a host object and save it. For details on creating TQLs, see "Creating a TQL Query" on page 59.

2 On the Report Manager, create a new report with the Report TQL you created in the previous step. For details on creating reports in the Report Manager, see "Creating a Report" on page 337.

3 After setting node definitions, right-click one of the report nodes in the Map pane and select **Node Order** to open the Node Order dialog box. For details, see "Setting the Node Order" on page 342.

4 Verify that the host node is at the top of the list or move it to this place by using the up arrow on the right.

5 Save the report by clicking the **Save** button on the toolbar.

To display a host report:

1 On the Map View, select the host for which you want to generate a Host Report from the Explorer, Map, or Information pane and right-click it, and select **Host Reports** to open the Reports on Host dialog box.
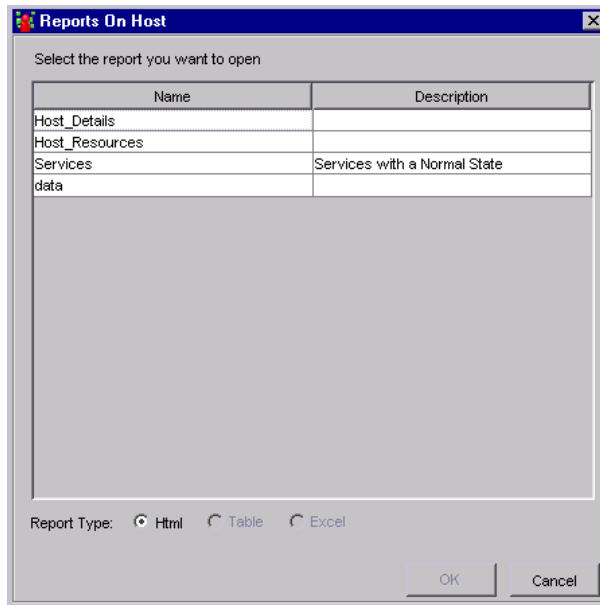
2 Select the Host Report you want to generate from the report list and click **OK**. A confirmation message is displayed:

3 Click **OK** to open the Report Download dialog box. For details, see step 7 in "Viewing and Saving the Generated Report" on page 345.

4 Select one of the following:

➤ To display the Host Report, select **Open this report** and click **OK**. The Host Report is displayed.

Produced By: **Mercury**   Creation Date **04/05/2005**

DNS Name	10.168.11.1
availability	
availabilitylast	0%

ip

Adress	Domain	Net Class	Net Address
10.168.11.1	niceDomain	C	10.168.11.0
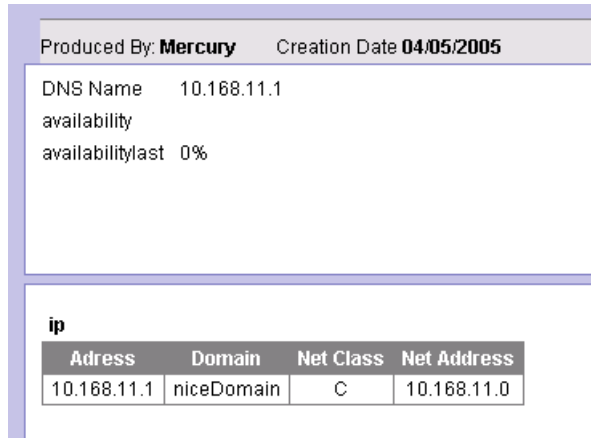
➤ To save the Host Report for future use, select the **Save this report on disk** option button on the Report Download dialog box, and click **OK**. A Save dialog box is displayed. Enter a unique name for the report and save it in your required location.

Part VI

Object Relationships

16

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 Explorer pane, enables you to define correlation rules, which specify a casual relationship between two or more objects – nodes and physical links. 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 objects. You can define the type and format of correlation events that are generated by root cause objects, how they change the Correlation and Admin States of the affected objects, and who is notified of the correlation occurrence.

A change in the Correlation State of an affected object can alter the object's Map Status, which is manifested through a change in the color of the object'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 would not blink, unlike all other cases in which the object contains unacknowledged events. This is caused by a built-in utility, which automatically suppresses the blink of an affected object'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 objects, while you determine and prioritize your responses to system alerts.

For example:

Explorer pane   *Map pane*
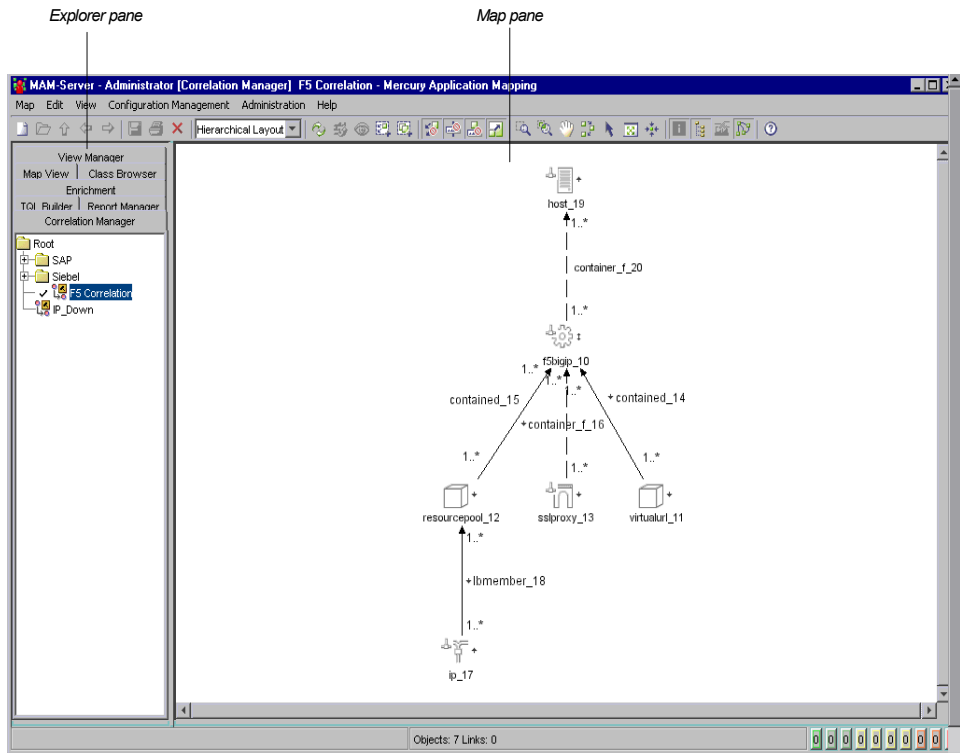
Correlation events are displayed in the Object Event, View Events, and All Event tabs in the Information pane of the Map View. An icon appears in the **iscorr** column, when the event is a result of a correlation occurrence. The relationship between root cause objects and affected objects can be examined in the Map View via the **Show Impact** and **Show Root Cause** options, as described "Showing a Correlation Impact" on page 189 and "Showing Root Cause Objects" on page 191.

When the Correlation Manager tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **Explorer pane –** 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.

➤ **Map pane –** Displays the currently selected rule, which consists of objects that are defined in the TQL query and the relationships between them. The trigger object is marked by an up arrow that is located to the right of the object's icon. Affected objects are marked by a down arrow to the right of the affected object's icon.

Correlation Manager Shortcut Menus

The Correlation Manager contains different right-click shortcut menus, depending on your selection.

This section includes:

➤

➤

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 Explorer pane:

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 database.
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.
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.

Option	Description
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.

Object/Link Shortcut Menu

The following table provides a brief description of each option in the shortcut menu displayed by right-clicking an object or a link in the Map pane:

Option	Description
Straight	(Enabled only for links with angles on the Map pane.) Enables you to straighten links with angles.
Define Affected	Displays the Root Cause dialog box, enabling you to define the affected object(s) and the correlation events that are generated by events in the root cause object.
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 objects are affected by the root cause object.

➤ Define the correlation events that are generated by events occurring in the root cause object.

➤ Define the format of messages sent to objects and users regarding correlation events.

➤ Show or hide selected nodes and links in the Map pane.

➤ 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 Map pane area to view the selected layer at different levels of magnification.

➤ Select different layout views.

17

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." You then proceed to define which object in this scenario is the root-cause object, and what objects are affected by it. This includes defining the actions that are performed on these affected objects, 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 objects and users are affected (for details, see "Defining Affected Objects 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 link definitions (for details, see "Defining Link Conditions" on page 72).

➤ 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 77 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 Explorer pane, selecting **Properties** from the shortcut menu, and selecting **Active**.

---

6 Click **OK**. The new correlation rule is displayed in the tree in the Explorer pane on the left side of the main window and in the Map pane on the right side of the window. For details of the Mercury Application Mapping main window, see "Mercury Application Mapping Main Window" on page 20.

Defining Affected Objects and Users

The next step in the procedure is to define which object in the TQL query is the root-cause object, that is, the trigger of the correlation event, and which objects 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 object

➤ defining how many instances of the root-cause object should fulfill the conditions for activating the correlation actions

➤ defining the objects that are affected by the root-cause object 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 objects to function as correlation triggers, they should comply with the following restrictions:

➤ You can select more than one object as a trigger. However, you cannot define an object as affected and as a trigger.

➤ If an object has a link whose minimum limit is **0** (meaning that one of its ends does not necessarily have an object linked to it), the object that is linked to its other end cannot be a root-cause object (since it may or may not exist in the TQL). For details about minimum limits, see "Defining Link Conditions" on page 72.

For example, object_15 cannot be either a root-cause or affected object because it is connected to the host with a Min limit of 0.


```
            host_14
        1..*      0..*
contained_17      container_f_18

        1..*        1..*
      ip_16       object_15
```
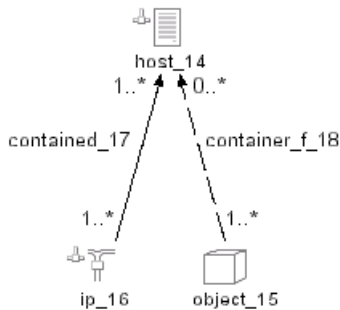
A node that is not visible cannot be a root-cause or an affected object.

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 objects.

This section includes the following topics:

➤ "Setting Multiple Conditions for Defining Affected Objects" on page 287

➤ "Editing a Correlation Rule" on page 291

➤ "Removing a Correlation Rule" on page 291

Setting Multiple Conditions for Defining Affected Objects

You can set multiple conditions for defining affected objects.

To set multiple conditions for defining affected objects:

1 In the Correlation Manager, right-click the object or link that functions as a trigger.

2 Select **Define Affected** to display the affected nodes.
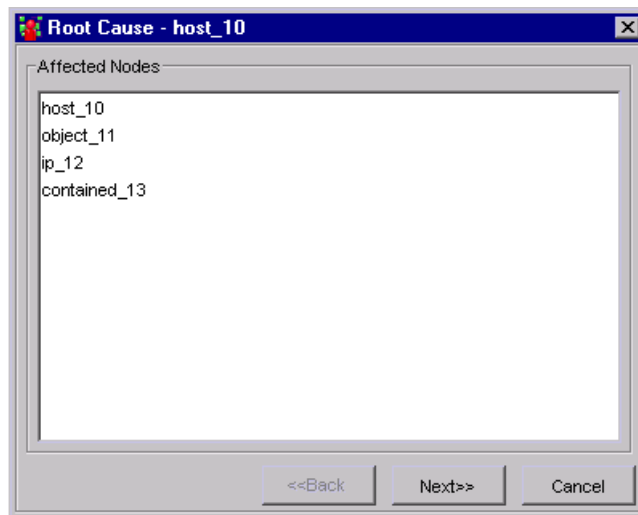
3 Select the node or nodes to be affected by the trigger.

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 object'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 object is different from Normal. When the Operational State of the object 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, Class 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 "Defining Node and Link Attribute Conditions" on page 70.

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 object 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 object) 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, as described in "Adding Attributes to the Label of the Selected Object" on page 101.

- 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 object 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 object.

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 object 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 Objects" 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 Map pane.

2 Click the **Save** button on the toolbar or select a different rule or tab in the Explorer pane. 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 74.

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 Explorer pane.

> **Note:** By default, an imported correlation rule is not active. To activate it, select it on the Explorer pane and right-click. From the shortcut menu, select **Properties**. In the displayed Properties dialog box, select **Active**, and click **OK**.

Part VII

Logical Objects and Rules

18

Introduction to Logical Objects and Links

This chapter introduces the Enrichment tab, which enables you to create logical objects and rules.

About Logical Objects, Links and Rules

The Enrichment tab 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 system.

Logical objects and links differ from other Mercury Application Mapping objects and links in that they are not actually discovered by discovery patterns. They are conceptual deductions that represent real links and objects that cannot be discovered automatically by the discovery system.

Logical objects and links are created as part of a TQL, whose other objects are regular ones, meaning objects that already exist in the database, as the following example illustrates:
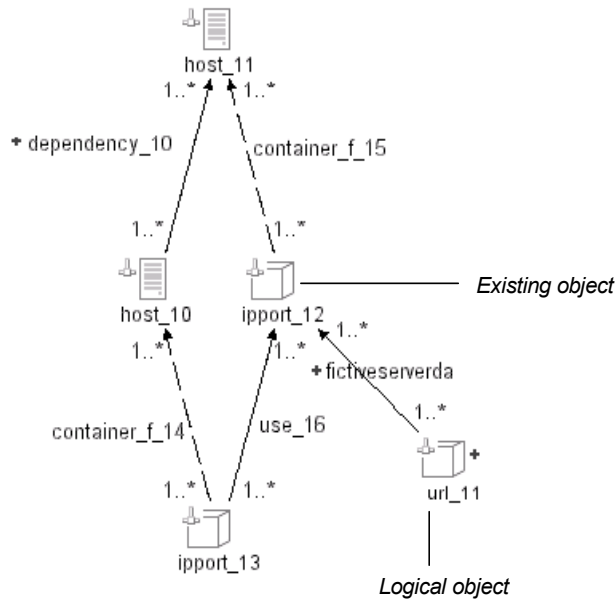
By placing a logical object within a TQL context, it receives data from the attributes of the other objects in the TQL, and uses it to insert new information into the database that is inferred from an existing one, or represents existing information in a new way.

Since these objects and links are usually characterized differently, their creation often entails a creation of new classes (done through the Class Browser; for details, see "Introduction to the Class Browser" on page 365).

Logical Object and Link Objectives

There are two main reasons for creating logical objects and links:

➤ Enlarging the Mercury Universal CMDB, by adding objects and links that are currently not included in it, but whose data is known or can be logically deduced from discovered objects.

For example, suppose it is known that a certain link, such as **depend**, exists between A and B, and between B and C, as shown in the following figure:
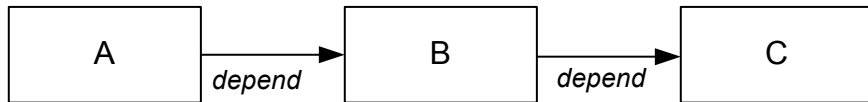
Based on this information, it is logical to deduce that between A and C a **depend** link also exists, as shown in the following figure:
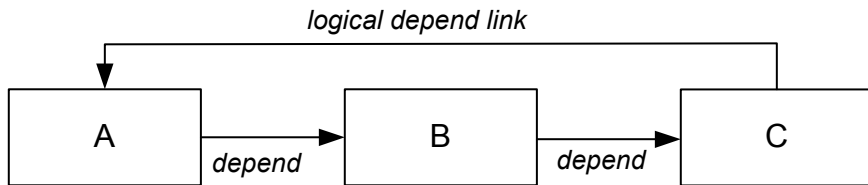
By using a logical link, you can add the link between A and C to the CMDB.

➤ Creating informative representations of your infrastructure, by simplifying data representations that are based on query results of complex structures.

Quick Tour of the Enrichment Tab

The **Enrichment** tab in the Explorer pane, enables you to create and define logical object rules. Theses rules, which are based on a specific Enrichment TQL query, can be used for several purposes: creating new logical objects and links for adding data to the CMDB, enabling additional representation options of existing data, updating the value of object attributes in the database, and deleting objects from the CMDB by using TQL queries.
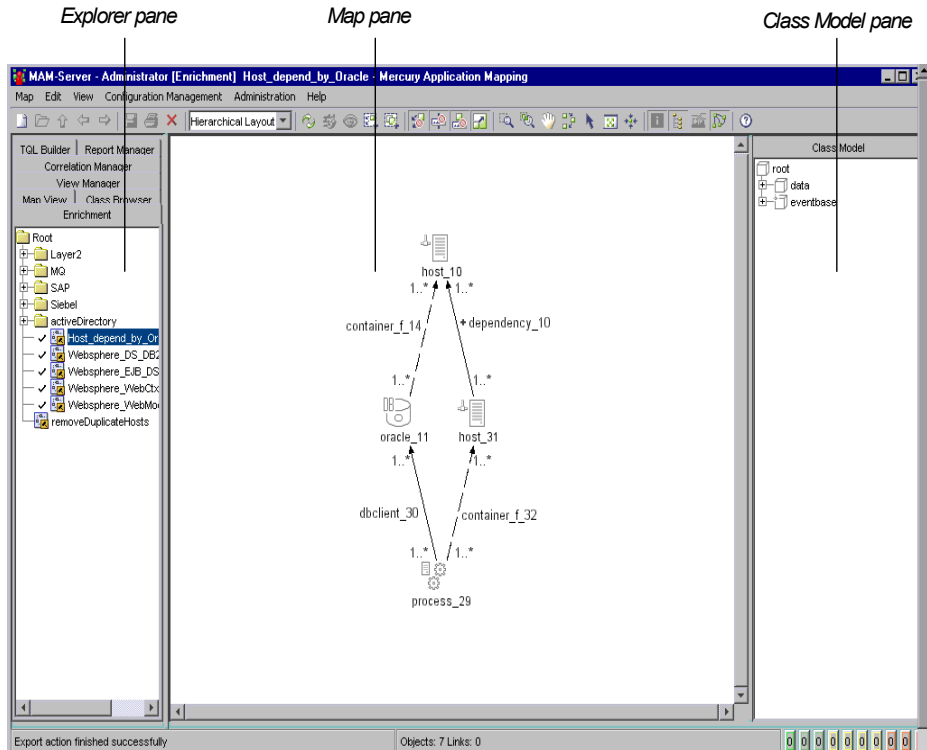
When the **Enrichment** tab is selected, Mercury Application Mapping main window is divided as follows:

➤ **Explorer pane –** Displays a hierarchical tree structure of the logical object rules you have defined. If required, you can export logical object rules and save them as XML scripts, which can be used for backup purposes.

Note: Each Logical Object Rule type is represented by a unique icon.

➤ **Map pane –** Displays the currently selected rule, which consists of objects that are defined in the TQL query and the relationships between them, and the logical objects and links that are created and added to the rule.

➤ **Class Model pane –** Represents the class model and contains icons for each object as defined by the administrator. By clicking and dragging objects to the Map pane, you can create new logical objects. By defining the relationship between these objects and existing ones, you can create new logical links.

Logical Object Builder Shortcut Menus

This section includes:

Logical Object Rule Shortcut Menu

The following table contains a brief description of each option in the shortcut menu displayed by right-clicking a selected Logical rule in the Explorer pane:

Option	Description
New	Enables you to define a new logical object rule.
New Folder	Enables you to create a new logical object folder.
Save	Enables you to save the logical rule and object to the database. (Enabled only when a new rule is created or when changes are made to an existing one.)

Option	Description
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).
Delete	Enables you to delete the selected rule from the database.
Export	Displays a standard Save As dialog box, enabling you to save the rule as an XML script. This option can be used to move selected rules from one workstation to another, provided the related TQL query is also relocated.
Properties	Enables you to change the properties of the rule (aside from the rule name and the Logical Object TQL query to which it is attached).

Logical Object/Link Shortcut Menu

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

Option	Description
Delete	Enables you to delete the selected logical object/link from the rule.
Node Definition	Displays the Node Definition dialog box, enabling you to define the attribute values of the logical object/link.

Regular Object Shortcut Menu

Option	Description
Delete Object/Link	Enables you to delete the selected link from the database.
Update Object/Link	Enables you to update the attribute values of the selected object/links.
Add Link	Enables you to add a logical link to the rule.

Working with the Logical Object Builder Manager

Use the Logical Object Builder to perform the following functions:

➤ Create new logical object rules, as well as duplicate and delete existing rules

➤ Define new logical objects and links and add them to the database

➤ Update the attribute values of existing objects

➤ Delete objects from the database

➤ Export and import selected logical object rules from XML scripts. This can be used to relocate logical object rules from one workstation to another. (This can be performed providing the related TQL query is also relocated.)

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

➤ Select different layout views.

19

Defining Logical Objects and Rules

This chapter explains how to define logical objects and rules.

About Defining Logical Objects and Rules

When creating a Logical Object rule, you first name the rule and attach it to a specific logical object TQL query. You then proceed to create new logical objects and links, by dragging and dropping the objects from the class model to the scenario on the Map pane, and connecting them to existing objects by logical and virtual links. You then set the attribute values of the logical objects and links for identifying the new objects and links in Mercury Application Mapping environment, and providing them with meaningful information.

Logical Object Rule Workflow

Logical object rules are created according to the following workflow:

➤ Create a logical object rule. For details, see "Creating a Logical Object Rule" on page 304.

➤ Create logical objects and links. For details, see "Creating Logical Objects and Links" on page 306.

➤ Define the logical objects and links. For details, see "Defining Logical Object and Link Attributes" on page 308.

➤ Save the logical object rule. For details, see "Saving the Logical Object Rule" on page 309.

Creating a Logical Object Rule

When creating a new logical object rule, you first need to perform the following:

➤ Define a unique name and a description for the rule

➤ Attach the rule to a specific Logical Object TQL

➤ Define whether the rule should be active in the system from the moment it is saved

➤ Determine under which conditions the new logical object(s) are deleted from the database

Bear in mind that any major changes made to the TQL query after creating a logical rule causes the system to delete the rule. These changes include the following:

➤ Delete, from the TQL, the object or objects connected to the logical object

➤ Change the **Min** and **Max** link definitions. For details, see "Defining Link Conditions" on page 72.

Minor changes, such as adding a node to the TQL, do not cause the deletion of the rule.

Note: The Logical Object TQL that serves as the basis for the logical 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 a logical rule.

To create a logical object rule:

1 On Enrichment, click **New** on the toolbar to open the New dialog box.

2 In the **Name** box, enter a unique name for the logical object rule.

3 (Optional) In the **Description** box, enter a description of the rule.

4 In the **Attached TQL** list, select the Logical Object TQL to which the logical object rule should be attached. For details on creating TQL queries, see Chapter 4, "Defining Topology Query Language (TQL) Queries."

5 (Optional) Verify that **Active** is selected, 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 Explorer pane and right-click the rule. From the shortcut menu, select **Properties**. In the Properties dialog box, select **Active**.

---

6 (Optional) To delete the logical objects that are created in this rule from the database once the rule is not active, clear **Active**, and select **Delete logical objects when the rule is not active**.

7 (Optional) To delete the logical object from the database once this rule is deleted, select **Delete logical objects when the rule is deleted**.

8 Click **OK**. The logical object rule is displayed in the Explorer pane on the left side of the main window and in the Map pane on the right side of the window.

Creating Logical Objects and Links

The next step in the procedure is to create logical objects and to define logical or virtual links between them and existing objects.

To create logical objects and links:

1 From the tree in the Explorer pane, select the logical object rule to which you want to add logical objects and links.

2 From the tree displayed in the class model, click and drag the object you want to function as a logical object to the Map pane. Added objects are displayed by an added indicator next to the object/link label.

Note: You can add more than one logical object to a rule.

---

3 The next step is to link the logical object to an existing object or objects, to provide the object with the needed context for its operation.

Note the following logical rule validations:

➤ You must link the new logical object to at least one of the existing objects in the rule.

➤ You cannot link one logical object to another.

➤ If the new logical object must be contained (according to its class definition) in another object, you must use the **container_f** link to connect this object to an existing one.

➤ You cannot link a logical object to a node that is not visible.

4 Select the logical object and the existing object between which you want to create a link, by holding down **Ctrl** and clicking the objects.

5 When both objects are selected, right-click to display a shortcut menu with two types of links for connecting the objects. Select one of the following:

➤ **Add Link –** Adds a regular link, which causes the creation of a link between the two objects in the CMDB. This option is similar to adding a link to a TQL in the TQL Builder. For details, see "Adding Nodes and Links to a Query" on page 66. When selecting this option, select the link type that defines the relationship between the two objects from the Add Link dialog box.

Note: The links that appear in the Add Link dialog box are the ones that are defined in the Class Browser as the relationship between the two object classes. If you create a new class for the logical object, and the link you need does not appear in the Add Link dialog box, return to the Class Browser and add a relationship between the two classes, as described in "Adding Relationship Links Between Classes" on page 386.

➤ **Add Virtual Link –** Adding a virtual link does not cause the creation of a real link in the Mercury Universal CMDB. However, you can use it to pass on information, such as an attribute value, from an existing object to a logical one.

✚ **Note:** Added links are displayed by an added indicator next to the object/link label.

Defining Logical Object and Link Attributes

After you have created the logical objects and links, you define their specific attribute values.

Logical Rule Validation

You must fill in the value of the key attribute(s) of the logical object. The method you use to define these values determines the number of instances created:

➤ If you enter a fixed value, you can create one instance only.

➤ If you enter a dynamic value, you can create numerous instances.

To define the logical objects and links:

1 In the Map pane, right-click the object or link whose attribute value you want to define, and select **TQL Node Definition** to open the Node Definition dialog box that displays the logical object or link attributes.

The Node Definition dialog box contains the following information about each attribute:

➤ **Name** – The name of the attribute.

➤ **Operator** – The type of operator.

➤ **Type** – The type of attribute (integer, string, and so on).

➤ **Value** – The value assigned to the attribute.

Note: Only attributes that have values or are marked as **Editable** in the Class Browser are displayed in the Node Definition dialog box. For details, see step 9 in "Creating Classes" on page 374.

---

2 To enter values for the object's attributes, click the **Edit** button. The object key attributes, which you must fill in, are displayed in bold. Depending on the attribute type you selected, one of the following Set Attributes dialog boxes is displayed:

➤ For an attribute of Boolean, Date, Integer or Enumeration type, enter a fixed value, or select a dynamic attribute value from the list and click **OK**. The attributes that appear in the **Attribute Name** list are the attributes of the connected regular object of the same type.

➤ For an attribute of String type, select one or more attributes from the list (for details, see "Adding Attributes to the Label of the Selected Object" on page 101), or enter a fixed value in the **Format** box, and click **OK**.

3 Click **OK**.

To delete an attribute value:

From the Node Definition dialog box, select the attribute you want to delete and click the **Clear** button or press the **Delete** key.

Saving the Logical Object Rule

After you have defined the logical object rule, the last step is to save it to the database. If you exit the Logical Object Builder without saving, a confirmation message is displayed, requesting that you save your rule.

To save the logical object rule:

Display the rule to be saved in the Map pane, and click **Save** on the toolbar or select a different rule or tab in the Explorer pane. A message is displayed, asking if you want to save the current logical object rule. Click **OK** to confirm the saving. The logical object 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 "Logical Rule Validation" on page 308.

Importing a Logical Object Rule

You can import XML files that contain saved logical object rules to your Logical Object Builder. Use this option if you want to relocate logical object rules from one workstation to another.

Note: Before you import a logical object rule, you must verify that its attached TQL query appears in the **Attached TQL** list on the New dialog box. For details, see "Creating a Logical Object Rule" on page 304. If the attached query does not exist in your query list, importing fails.

By default, an imported logical object rule is not active. This section explains how to activate an imported logical object.

To import a logical object rule:

1 From the Logical Object Builder, open the **Map** menu and click **Import** to open the Import dialog box. For details about this dialog box, see "Importing a TQL Query" on page 74.

2 Locate the logical object rule you want to import, select it and click **Import**. The imported logical object rule is added to the logical object rule list in the Explorer pane.

To activate an imported logical object rule:

1 Right-click the rule you want to activate on the Explorer pane and select **Properties**.

2 In the Properties dialog box, select **Active**.

3 Click **OK**.

Updating Object Attributes Using Logical Object Rules

Use a logical object rule to update the value of object attributes in the database, or to add data to attributes that currently do not have values (for details, see "Creating a Logical Object Rule" on page 304). You can use this option, for example, for adding a note simultaneously to all instances of an object.

When updating object attributes, you do not have to create logical objects and links. You open a Logical Object TQL query in the Logical Object Builder, and update the attributes of one or more of its objects.

To update object attributes:

1 On the Logical Object Builder, click **New** on the toolbar to open the New dialog box.

2 In the **Name** box, enter a unique name for the logical object rule.

3 Enter a description of the object in the **Description** box.

4 From the **Attached TQL** list, select the Logical Object TQL to which the logical object rule should be attached, and click **OK**.

5 On the Map pane, right-click the object whose attribute(s) you want to update and select **Update Object/Link** to open the Node Definition dialog box. For details, see "Defining Logical Object and Link Attributes" on page 308.

6 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. For details, see step 2 in "Defining Logical Object and Link Attributes" on page 308.

7 Enter the updated value and click **OK**. Updated objects are displayed by an arrow indicator next to the object/link label.

Note: If you want to cancel the update and you have not yet saved the rule, right-click the updated object and select **Clear**. If you already saved the rule, you cannot reverse the update.

---

8 To save the rule and update the objects, click **Save** on the toolbar.

Deleting Objects from the Mercury Universal CMDB Using Logical Object Rules

The Logical Object Builder enables you to delete objects from the Mercury Universal CMDB by using TQL queries. You create a specific logical object rule that is designed for an object deletion. You can use this option, for example, for removing unnecessary data from the database, or for automatically deleting copies of, or partial, objects.

Note: To delete logical objects and links, use the deletion options that appear in the rule's New or Properties dialog box. For details, see"Creating a Logical Object Rule" on page 304.

To delete an object from the database:

1 On the Map pane, right-click the object whose instances you want to delete from the database and select **Delete Object/Link**. Deleted objects are displayed by a deleted indicator next to the object/link label.

Note: To cancel the deletion if you have not yet saved the rule, right-click the deleted object and select **Clear**. If you have saved the rule, you cannot reverse the deletion.

---

2 To save the rule and update the objects, click **Save** on the toolbar.

20

Creating a Sample View with Logical Objects

This chapter introduces the Logical Object Builder Manager, which enables you to create Logical Objects and Rules.

This chapter describes:	On page:
Activating Discovery Patterns	314
Creating a Logical Object TQL Query	315
Creating a Class for a New Logical Object	318
Creating and Defining the Logical Object and Rule	320
Displaying the Logical Object View	324

About Creating a Sample View

The following section leads you step-by-step through the process of creating a sample view with a logical object. The purpose of the sample view is to discover and display the location of your network's hosts. To generate such a view, which visualizes your network's physical deployment and reveals the requested information, you can use several discovered objects together with a logical one.

Activating Discovery Patterns

Location details are often entered as an SNMP agent's attribute in the installation process. To use this data in a way that reveals the host locations, you need to establish a connection between hosts and their switches, which contain the SNMP agents, assuming that they are located at the same place. Mercury Application Mapping provides you with the necessary discovery patterns for discovering the following objects and their links, which establish the requested connection:

➤ Host, interface, and their link

➤ Port and its link to an interface

➤ Port and its link to a switch

➤ Switch and its link to its SNMP agent

Mercury Application Mapping enables administrators to activate or deactivate selected discovery patterns from among the list provided in the system.

An activated discovery pattern invokes one or more discovery tasks, which triggers discovery methods to begin discovering the data defined in the pattern.

For example, if you select **SNMP_NET_Dis_Bridge**, the discovery methods begin to discover all the bridges on a target host (followed by their ports and Layer2 connections), which comply with the definition contained in that specific TQL. The administrator activates the required patterns before users begin working with the system.

To activate discovery patterns:

1 Open the **Administration** menu and select **Discovery Pattern Manager** to open the Discovery Pattern Manager dialog box.

2 In the **Packages** tab, expand the Network package and select following discovery patterns: **ICMP_NET_DIS_IpC**, **SNMP_NET_Dis_Connection**, **SNMP_NET_Dis_HostBase**, and **SNMP_NET_Dis_Bridge** patterns. Selecting those patterns in the **Packages** tab causes them to be selected everywhere they appear in the **Packages** and **Modules** tabs.

3 Click **Apply** to activate them.

Note: If you use switches of one of the following types: Enterasys3, Enterasys5, or Enterasys Vlan, activate the patterns that carry these names (such as **SNMP_NET_Dis_EndU_Enterasys3_Dynamic**) instead of the **SNMP_NET_Dis_Bridge** pattern.

---

4 Click **OK** to activate the selected discovery patterns and to close the dialog box.

Creating a Logical Object TQL Query

The appropriate discovery patterns can establish a connection between hosts, their switches, and SNMP agents. You must create a Logical Object TQL as a basis for the logical object rule that uses this connection to discover and display the host locations.

To create a Logical Object TQL query:

1 On the TQL Builder, click the **New** button to open the Create New TQL dialog box (for details, see "Creating a TQL Query" on page 59).

2 Enter the following information in the Create New TQL dialog box:

> ➤ In the **TQL Name** box, enter **Hosts**.

> ➤ From the **Type** list, select **Enrichment**.

> ➤ From the **Priority** list, select **Low**.

3 Click **OK**. The Hosts TQL is displayed in the Explorer pane on the left side of the main window.

4 With the Hosts TQL selected, click and drag the following objects to the Map pane from the class model on the right side of the main window:

> ➤ Host

> ➤ Interface

> ➤ Port

➤ Switch

➤ SNMP

5 Select the host and interface objects and right-click.

➤ Select **Add Link** to open the Add Links dialog box.

➤ Select the **interface** > **container_f** > **host link**, and click **Finish**. The first pair of host and interface objects is connected by the container_f link.

➤ Repeat this procedure to define the other links as follows:

- interface–port: **layertwo**

- port–switch: **container_f**

- switch–snmp: **container_f**

The Logical Object TQL now looks similar to the following example. The example shows the objects and links that are discovered by the discovery patterns and the way they are interconnected:
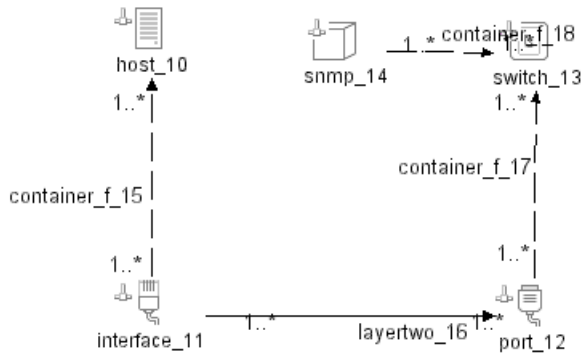
6 Since only SNMP agents (and their hosts) that have location data should be included in the view, you can add a condition to the SNMP object that excludes SNMP agents without location data:

➤ Right-click the SNMP object, and select **TQL Node Definition** to open the TQL Node Definition dialog box. For details, see step 1 in "Defining Node and Link Attribute Conditions" on page 70.

➤ On the Attribute Condition tab, click **Add** to open the Condition dialog box. For details, see step 2 in "Defining Node and Link Attribute Conditions" on page 70.

➤ From the **Attribute Name** list, select the **snmp_location** attribute.

➤ From the **Operator** list, select the **IS NOT NULL** operator, and click **OK** to close the Condition dialog box. The TQL Node Definition dialog box now appears as follows:
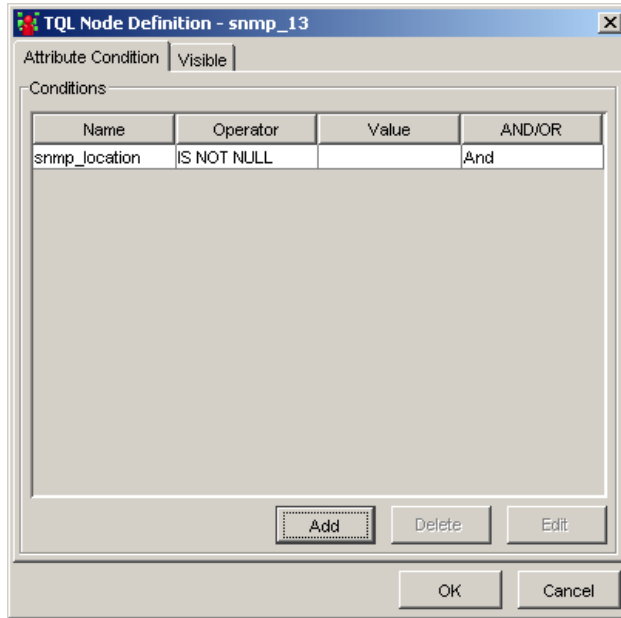
➤ Click **OK** to save the new condition and close the TQL Node Definition dialog box.

7 To save the Hosts TQL, click the **Save** button on the TQL Builder toolbar.

Creating a Class for a New Logical Object

The Location object that visualizes the host locations is not based on one of the existing classes in the system. Therefore, you must create a class for it.

To create a class:

1 On the Class Browser, click **New** to open the Class Wizard.

2 Enter **Location** in the **Class Name** box.

3 From the **Base Class** list, select the **Object** class as the base class of the class.

4 Click **Next** to open the Class Attributes dialog box.

5 Define the **data_name** attribute as the location class key attribute, by clicking in the left column, as shown in the following figure:
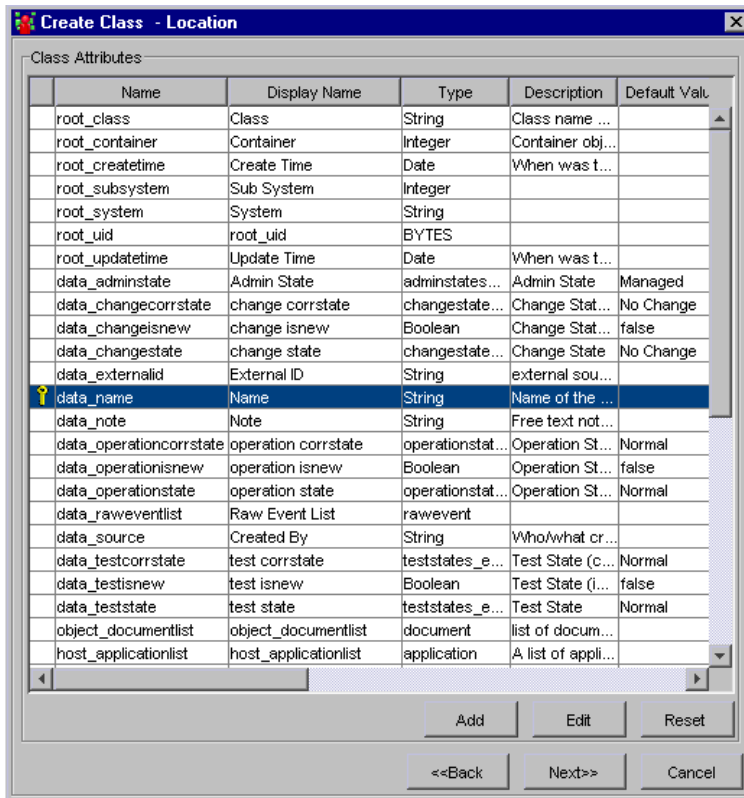
6 To add new attributes, click **Add** to open the Add Attribute dialog box.

For details on modifying field properties, see step 9 in "Creating Classes" on page 374.

7 Click **OK** to save the changes you have made. The class attribute appears in the Class Attributes dialog box.

8 Click **Next** to attach an icon to the class. For details, see "Assigning an Icon to a Class" on page 378.

9 Click **Next** to create a class method. For details, see "Creating Class Methods" on page 380.

10 Click **Next** to define new class tags. For details, see "Adding Tags to the Class" on page 384.

11 Accept the default selection and click **Next** to open the Class Default Label page.

12 Click **Finish**. The new Location class is created and displayed in the Explorer pane.

To create a real link between the location class and the host:

To create a real link between the location class and the host and switches classes, you must define the relationship between them, as follows:

1 On the Explorer pane, select the **Location** class, click CTRL, and select **host** classes and right-click.

2 Select **Add/Remove Relationship** to open the Add Relationship dialog box.

3 To add a member link, select the check box on the **host** > **Location** column, as shown in the following figure:
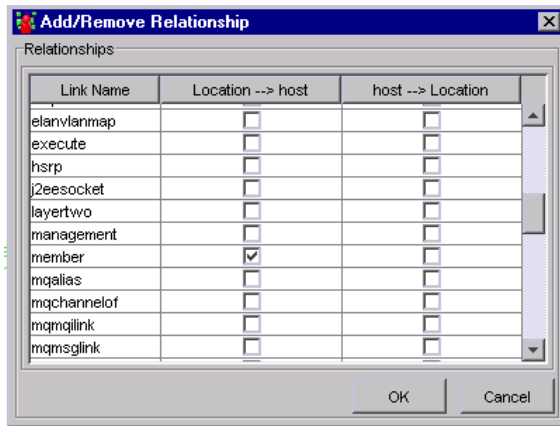
4 Click **OK**. A new member link is created between the location and host classes. Since the switch class inherits from the host class, a new member link is automatically created between the location and switch classes as well.

You can use the class for the creation of the new logical object.

Creating and Defining the Logical Object and Rule

After the creation of the TQL and the establishment of a connection between a host and a switch, you can add a logical object, called **Location**, and connect it to the host and switch.

To create and define the logical object and rule:

1 In the Enrichment tab, click the **New** button to open the New dialog box. For details, see "Creating a Logical Object Rule" on page 304.

2 In the **Name** box, enter **Host Location** as the name of the new logical rule.

3 From the **Attached TQL** list, select **hosts.**

4 Select **Active**.

5 Click **OK.**

The new logical rule is displayed in the Explorer pane on the left side of the main window, and in the Map pane in the middle of the window.

6 From the class model on the right, drag and drop the Location class in the Map pane, as shown here:
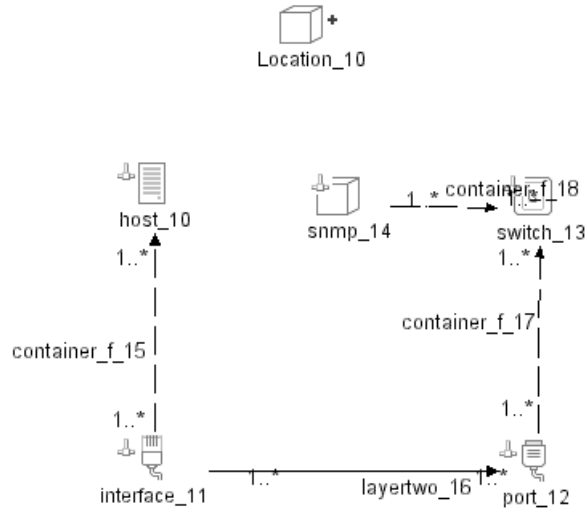
7 To establish a link between **host** and **Location**:

> ➤ Select the objects and right-click.

> ➤ Select **Add Link** to open the Add Links dialog box.

> ➤ Select **host** > **member** > **Location**

> ➤ Click **Finish**.

8 Repeat the previous step to establish a **member** link between the **switch** and **Location** objects.

The following scenario is created:
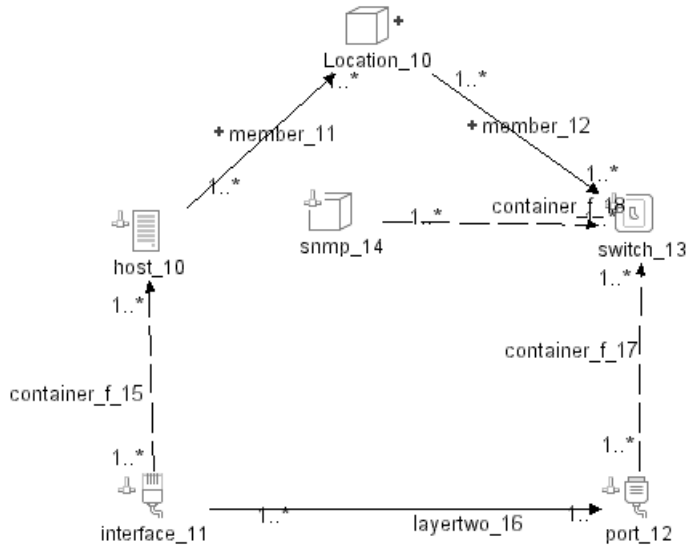
9 To define the key attribute of the Location logical object, select it and right-click.

➤ Select **Node Definition** to open the Node Definition dialog box. For details, see "Logical Rule Validation" on page 308.

➤ From the Name column, select the **Name** attribute and click **Edit** to open the Set Attributes dialog box. For details, see step 2 in "Logical Rule Validation" on page 308.

➤ Select the **snmp_14.snmp_location** attribute and click the plus button next to it.

Note: In your system, the **snmp_location** attribute may be labeled differently.

The selected attribute is displayed in the **Format** box, as shown in the following figure:
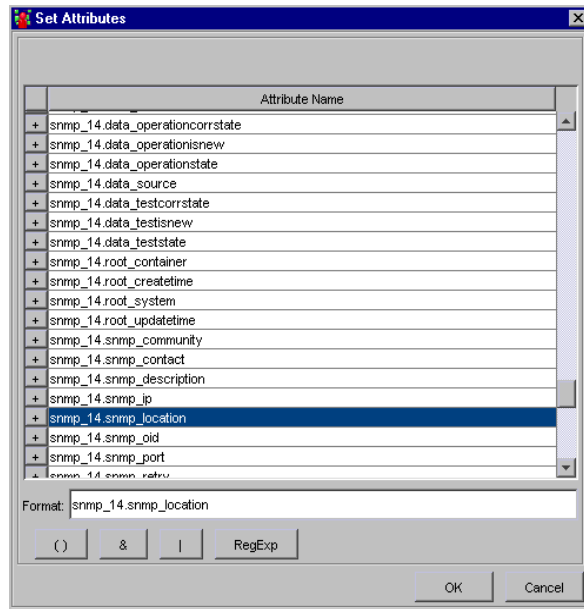
10 Click **OK**. You are returned to the Node Definition dialog box.

11 Click **OK** again to save your definition and close the dialog box.

12 To save the new logical rule, click the **Save** button on the toolbar.

The Location object is added to the CMDB and can be used in queries and views. You can use it to create a simple view that displays the location of your network's hosts. The information that is displayed in this view is deduced from the **snmp_ location** attribute.

Displaying the Logical Object View

To display the logical object which contains the host location data, you need to create a View TQL and then a new view for it.

To create a View TQL:

1 On the TQL Builder, click the **New** button to open the Create New TQL dialog box. For details, see "Creating a TQL Query" on page 59.

2 Enter the following information in the Create New TQL dialog box:

➤ In the **TQL Name** box, enter **Host Location**.

➤ From the **Type** list, select **View**.

3 Click **OK** to display the Host Location TQL in the Explorer pane.

4 With the Host Location TQL selected in the Explorer pane, click and drag to the Map pane, from the class model on the right, the following objects:

➤ Location

➤ Host

5 Create a member link between the Location and host objects, as follows:

➤ Select the **host** and **Location** objects and right-click.

➤ Select **Add Link** to open the Add Links dialog box.

➤ Select the **host** > **member** > **Location** link

➤ Click **Finish**. The objects are connected by the member link.

The Host Location TQL should look like this:
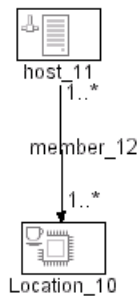
6 To save the Host Location TQL, click the **Save** button on the TQL Builder toolbar.

To create a view:

1 In the View Manager, click the **New** button on the toolbar to open the Create New View dialog box. For details, see "Creating a View" on page 92.

2 In the Create New View dialog box, enter the following information:

> ➤ In the **View Name** box, type **Host Location**.

> ➤ From the **Attached TQL** list, select **Host Location**.

> ➤ Select **Merge Identical Instances**.

3 Click **OK** to save definitions and close the Create New View dialog box.

4 To save the Host Location view, click the **Save** button on the toolbar.

5 Verify that a checkmark is displayed beside the Host Location view in the Explorer pane (indicates that the view is to be displayed in the Map View). If not, right-click the view in the Explorer pane and select **Add/Remove to Map View**.

You can now view the results of the logical object rule you created. On the Map View, select the **Host Location** view from the Explorer pane.

Part VIII

System Reports

21

Introduction to the Report Manager

This chapter introduces the Report Manager, which enables you to create reports.

This chapter describes:	On page:
Report Manager Shortcut Menus	331
Working with the Report Manager	333

About Introduction to the Report Manager

The Report Manager is displayed by selecting the Report Manager tab in the Explorer pane. It 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.

View Explorer   *Topology Map*
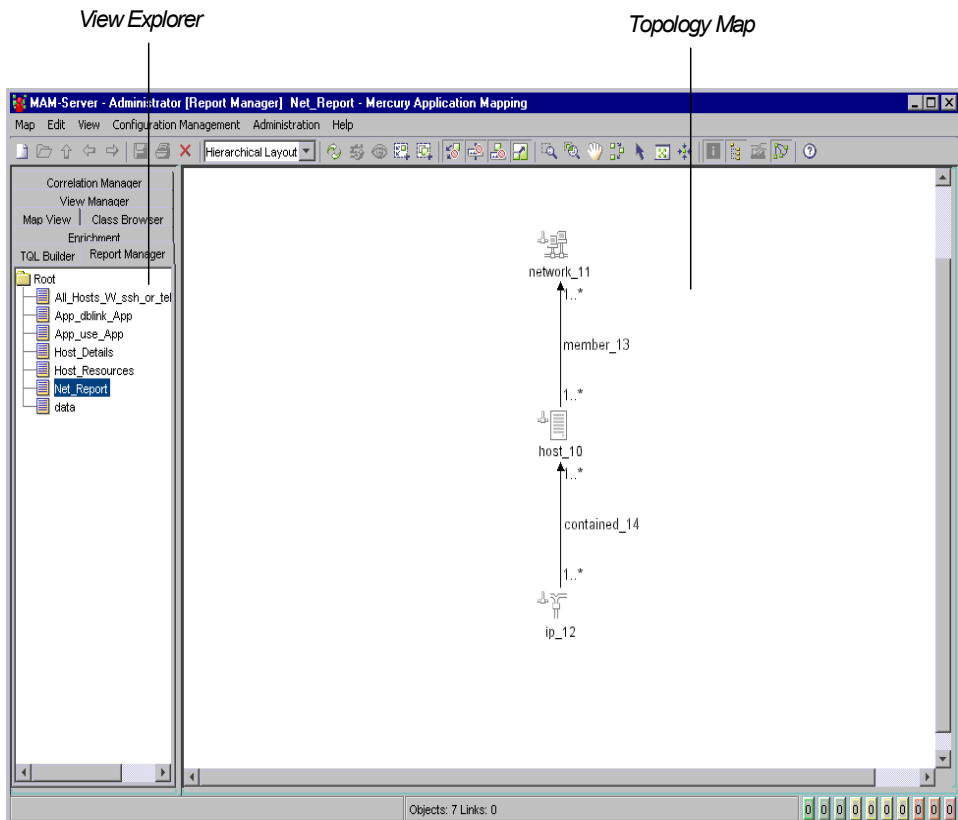
When the Report Manager tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **Explorer pane** – 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.

➤ **Map pane** – Displays the currently selected report, which consists of objects and the relationships between them. The displayed scenario is defined as a Report TQL query in the TQL Builder.

Report Manager Shortcut Menus

The Report Manager contains different shortcut menus, depending on your selection, as follows:

➤ "Report Shortcut Menu" on page 331

➤ "Object Shortcut Menu" on page 332

➤ "Link Shortcut Menu" on page 332

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 Explorer pane:

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).
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).

Object Shortcut Menu

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

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 node.
Node Order	Displays the Set Node Order dialog box, enabling you to determine the order in which node information is displayed in the report, that is, the order of the report's columns.

Link Shortcut Menu

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

Option	Description
Sibling	Defines one node as a sibling of the node to which it is linked. In the report, the two nodes appear side by side.
Child	Defines one node as a child of the node to which it is linked. In the report, the child node appears beneath (and indented from) the parent node.
Node Order	Displays the Set Node Order dialog box, enabling you to determine the order in which 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 nodes should appear in the report, the parameter reported for each node and the function to be performed on the information

➤ Define the sort order of the information reported for each node

➤ Determine the order of nodes in the report

➤ Define which nodes appear as children beneath parent nodes

➤ Create and delete tags, which are logical groupings of objects and links, to be defined in the query

➤ 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 Map pane area to view the selected layer at different levels of magnification

➤ Select different layout views

22

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 337).

➤ Define the report nodes (for details, see "Defining the Report Nodes" on page 337).

➤ Set the node order (for details, see "Setting the Node Order" on page 342).

➤ Refine the report layout (for details, see "Refining the Report Layout" on page 343).

➤ Save the report template (for details, see "Saving the System Report Template" on page 344).

➤ View and save the generated report (for details, see "Viewing and Saving the Generated Report" on page 345).

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.

 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 Explorer pane and in the Map pane.

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 Map pane, and select **Report Node Definition** to open the Report Node Definition dialog box.

 2 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.

 3 Click the **Add** button to display the Column Definition wizard.

   You define the node's columns that consist of the node attributes and functions.

 4 Select **Attribute value columns** and click **Next**.

 5 To add a column to the report relating to the selected node:

   ➤ In the **Attribute List** area, select an attribute and click the 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 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.

6 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.

For example, to display the availability percentage of the selected node, select **Availability**. (You can select availability on condition that you previously defined availability rules for this node, as described in "Adding Availability Rules" on page 115.)

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.

---

7 (Optional) If the selected function requires additional definitions, click **Next** to display a wizard and define the necessary parameters.

 8 Click **Finish** to close the Column Definition Wizard. The new column
 definitions are displayed in the Set Columns tab of the Node Definition
 dialog box, as shown in the following figure:
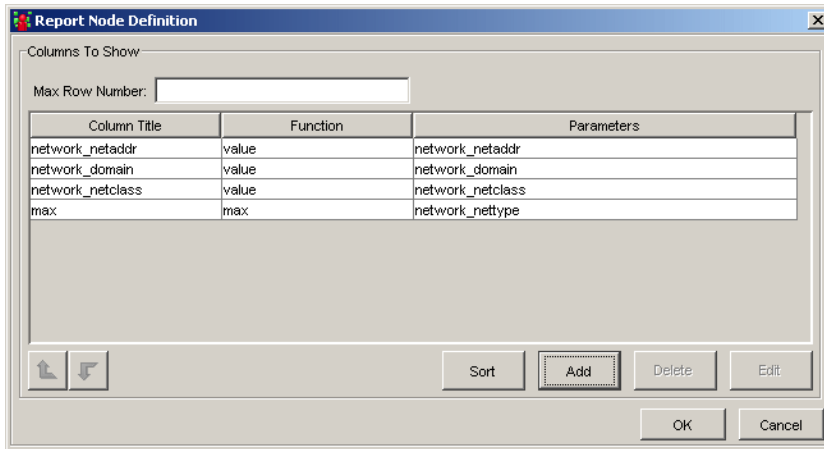
 Each row in the table represents a column that appears in the report.

 9 Repeat steps 3 to 5 to add additional column definitions.

 Note:

 ➤ You can change the order of the columns (for details, see "Setting the
 Node Order" on page 342).

 ➤ You can sort the information in the columns (for details, see "Sorting
 Column Information" on page 341).

---

 10 Click **OK** to close the Report Node Definition dialog box.

 11 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 Map pane, 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 Map pane, 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 Map pane, 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 Map pane 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):

network_netaddr	network_domain	network_netclass	host_label	host_key	ip_address	ip_netaddr	ip_domain
			HPSVR01	080009C40BC0	192.168.0.3	192.168.0.0	niceDomain
			LT-SCOTT-S	000A940030CA	192.168.0.22	192.168.0.0	niceDomain
			MAMPOC	00301BB66A55	192.168.0.6	192.168.0.0	niceDomain
			RH9SVR01	006097D87BBA	192.168.0.4	192.168.0.0	niceDomain
192.168.0.0	niceDomain	C	SCARECROW	00041F0520CC	192.168.0.1	192.168.0.0	niceDomain
			SUNSVR01	0800208EE861	192.168.0.2	192.168.0.0	niceDomain
			WINSVR01	00012921CC36	192.168.0.9	192.168.0.0	niceDomain
			WINSVR02	000D87BAD9A3	192.168.0.10	192.168.0.0	niceDomain
			WINSVR03	000D87B2A210	192.168.0.11	192.168.0.0	niceDomain

➤ The child node (IP) is displayed below the parent node (Host):

network_netaddr	network_domain	network_netclass			
192.168.0.0	niceDomain	C			
	host_label	host_key	ip_address	ip_netaddr	ip_domain
	HPSVR01	080009C40BC0	192.168.0.3	192.168.0.0	niceDomain
	LT-SCOTT-S	000A940030CA	192.168.0.22	192.168.0.0	niceDomain
	MAMPOC	00301BB66A55	192.168.0.6	192.168.0.0	niceDomain
	RH9SVR01	006097D87BBA	192.168.0.4	192.168.0.0	niceDomain
	SCARECROW	00041F0520CC	192.168.0.1	192.168.0.0	niceDomain
	SUNSVR01	0800208EE861	192.168.0.2	192.168.0.0	niceDomain
	WINSVR01	00012921CC36	192.168.0.9	192.168.0.0	niceDomain
	WINSVR02	000D87BAD9A3	192.168.0.10	192.168.0.0	niceDomain
	WINSVR03	000D87B2A210	192.168.0.11	192.168.0.0	niceDomain

To refine the report layout:

1 On the Map pane of the Report Manager, right-click a link 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 link 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 Map pane, and click **Save** on the toolbar or select a different report or tab in the Explorer pane. 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.

Net_Report								
network_ne...	network_do...	network_ne...	host_label	host_key	ip_address	ip_netaddr	ip_domain	
192.168.0.0	niceDomain	C	HPSVR01	080009C40B...	192.168.0.3	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	LT-SCOTT-S	000A940030...	192.168.0.22	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	MAMPOC	00301BB66...	192.168.0.6	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	RH9SVR01	006097D87B...	192.168.0.4	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	SCARECROW	00041F0520...	192.168.0.1	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	SUNSVR01	0800208EE861	192.168.0.2	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	WINSVR01	00012921CC...	192.168.0.9	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	WINSVR02	000D87BAD...	192.168.0.10	192.168.0.0	niceDomain	
192.168.0.0	niceDomain	C	WINSVR03	000D87B2A...	192.168.0.11	192.168.0.0	niceDomain	

To view and save the report:

 1 From the **Reports** menu, right-click the appropriate 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 347.

  ➤ **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 **Reports** > **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.
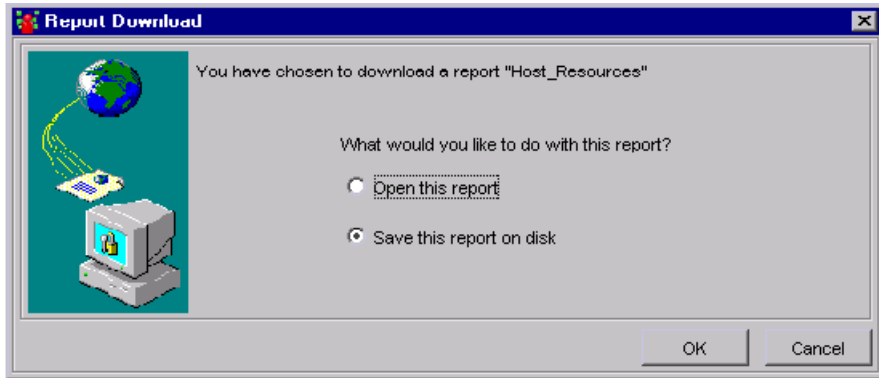
6 Click **Open this report** and click **OK** if you want to open the report or click **Save this report on disk** and click **OK** if you want to save the report. The Save dialog box opens, enter a unique name for the report and save it in the required location.

➤ 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:

network_netaddr	network_domain	network_netclass	host_label	host_key	ip_address	ip_netaddr	ip_domain
			HPSVR01	080009C40BC0	192.168.0.3	192.168.0.0	niceDomain
			LT-SCOTT-S	000A940030CA	192.168.0.22	192.168.0.0	niceDomain
			MAMPOC	00301BB66A55	192.168.0.6	192.168.0.0	niceDomain
			RH9SVR01	006097D87BBA	192.168.0.4	192.168.0.0	niceDomain
192.168.0.0	niceDomain	C	SCARECROW	00041F0520CC	192.168.0.1	192.168.0.0	niceDomain
			SUNSVR01	0800208EE861	192.168.0.2	192.168.0.0	niceDomain
			WINSVR01	00012921CC36	192.168.0.9	192.168.0.0	niceDomain
			WINSVR02	000D87BAD9A3	192.168.0.10	192.168.0.0	niceDomain
			WINSVR03	000D87B2A210	192.168.0.11	192.168.0.0	niceDomain

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 a 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 74).

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 Explorer pane.

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 345.

2 Right-click the table and select **Chart Manager** to open the Chart Editor dialog box. For details, see "Converting a Report into a Chart" on page 240.

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 241.

23

Creating a Sample System Report

This chapter explains how to create a sample system report.

This chapter describes:	On page:
Creating the Report TQL Query	350
Creating the Report	353
Defining the Report Nodes	356
Sorting the Node Data	358
Setting the Node Order	359
Refining and Saving the Report Layout	360
Viewing and Saving the Report	361
Converting the Report to a Chart	362

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 **New** button 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 Explorer pane on the left side of the main window, in the Report folder.

 7 With the Services TQL selected from the class model on the right side of the main window, click and drag the following objects to the Map pane:

➤ Network

➤ Host

➤ Service

The Map pane looks something like this:

8 Create a links between the network and host objects and between the host and service objects as follows:

➤ Select the **network** and **host** objects and right-click them.

➤ Select **Add Link** to open the Add Links dialog box.

➤ Select the **host** > **member** > **network** link, and click **Finish**. The first pair of host and network objects is connected by the member link.

➤ Repeat this process to define another container link between the service and host objects. For these objects, select the **service** > **container_f** > **host** link from the Add Link dialog box.

The Report TQL now displays as follows:
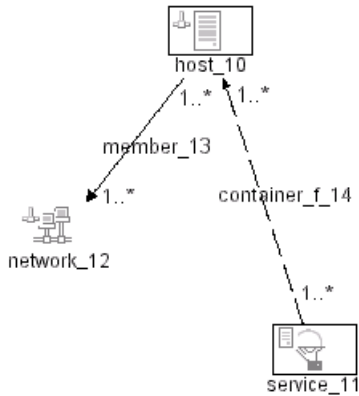
9 The next step is to define the Report TQL to retrieve only services that are indicated as malfunctioning. This is done by defining the attribute condition of the host object to retrieve only services that are in a state that is higher than the Normal state. Set this definition as follows:

➤ In the Map pane, right-click the **host** object 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 **Operator** list, select more than (>).

➤ In the Value area, select **Normal**.

➤ Click **OK**. The TQL Node Definition dialog box now appears as shown in the following figure:
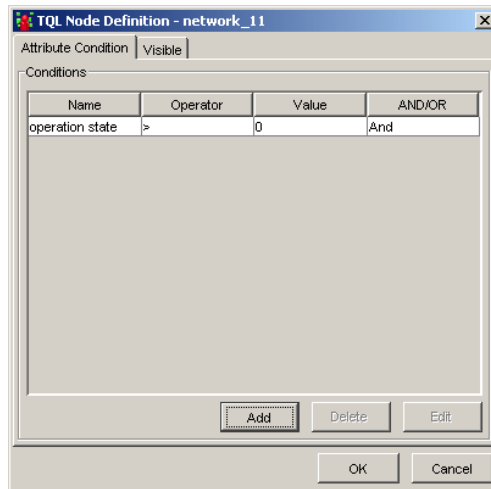
➤ Click **OK**. When Operation State is more than Normal, all services that are higher than Normal are displayed in the report query results.

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 **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:
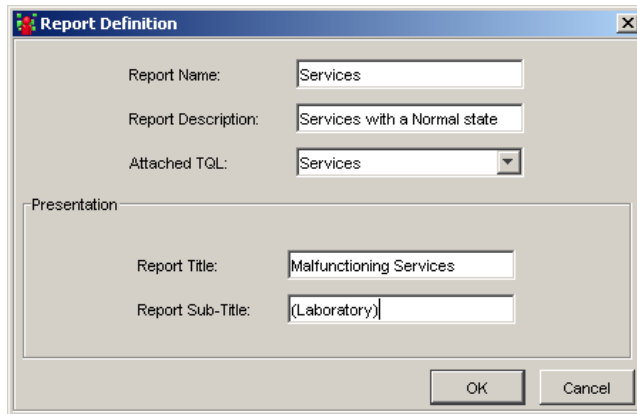
8 Click **OK**. The new report is displayed in the Explorer pane on the left side of the main window, and in the Map pane on the right side of the window, as shown in the following figure:
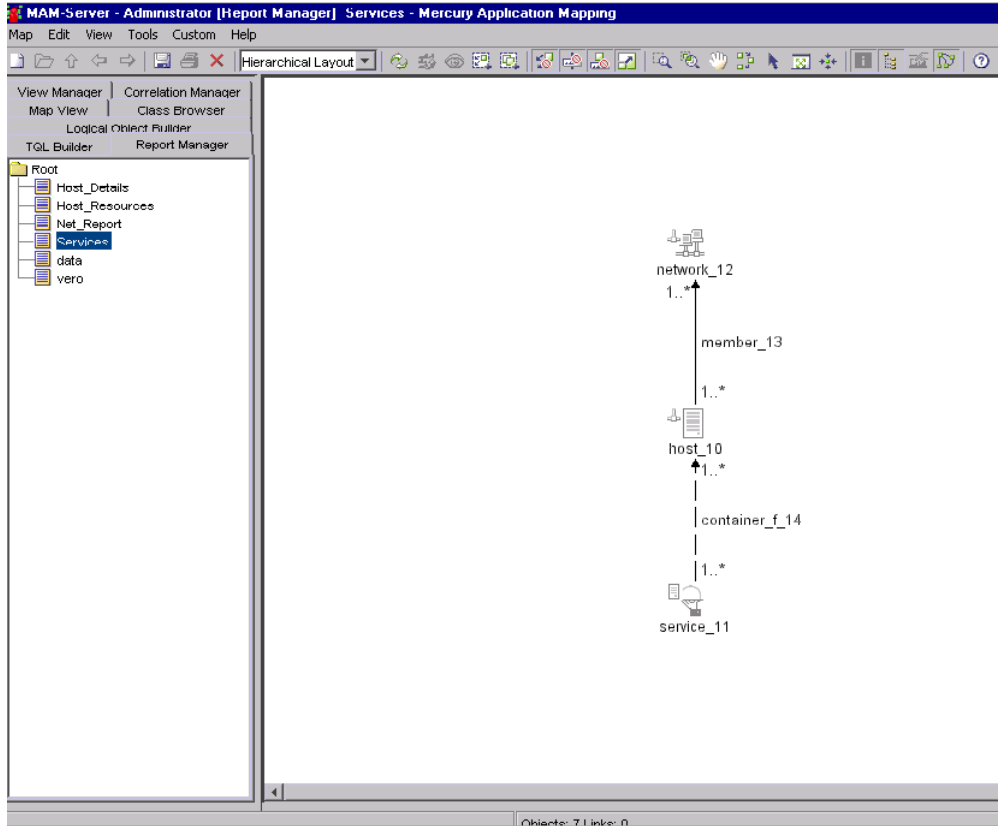
Defining the Report Nodes

After creating the report, the next step is to define the 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 node (attribute and function), the maximum number of rows to allow for each node and the node label (that is, the column title).

In this report you have three nodes: network, host, and service. For each 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 node, select each of the nodes in the scenario displayed in the Map pane.

1 With the node selected, right-click and select **Report Node Definition** to open the Report Node Definition dialog box.

2 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.

3 Click the **Add** button to display the Column Definition Wizard.

Define the number and content of the selected 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.

4 Select **Attribute value columns** and click **Next**.

5 In the **Attribute List** area, select the following attributes for the **network** node and click the arrow button:

➤ **network_domain**

➤ **network_netclass**

➤ **network_nettype**

➤ **operation state**

6 After you move a selected attribute, select it in the **Column List** area. Enter a new title in the **Column Title** box according to the titles specified on the Column Title column in the table above.

Node	Attribute Name	Column Title
Network	network_domain	network_domain
	network_netclass	network_netclass
	network_nettype	network_nettype
	Operation State	Operation State

7 Click **Apply**.

8 Click **Finish**. You are returned to the first page of the Column Definition Wizard.

The Report Node Definition dialog box of the network node now displays as follows:
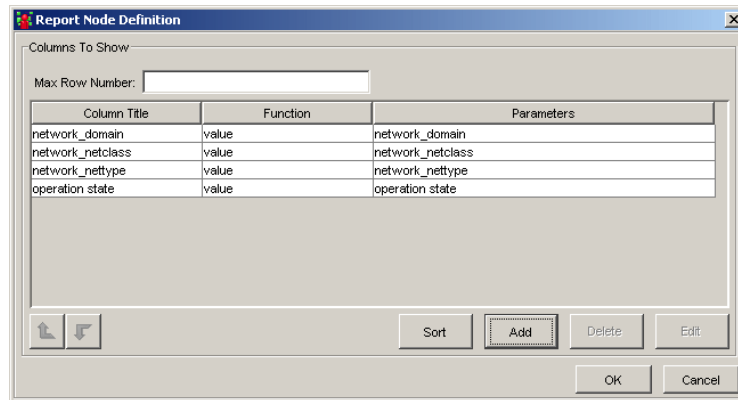
Each row in the table represents a column that will appear in the report.

9 To change the order of the network's columns, use the up and down arrows.

357

10 Repeat steps 1 to 6 to add column definitions for the host and service nodes. In the Column Definition Wizard, select and enter the following:

➤ For the **host** node:

- **host_dnsname**
- **Operation State**

➤ For the **service** node:

- **Name**
- **service_operatingstatus**
- **operation state**

Sorting the Node Data

You must now determine the order of the information displayed for each node, and the method by which the data is sorted in each column.

To sort the node data:

Perform the following steps for each node displayed in the Map pane.

1 Right-click the appropriate node and select **Report Node Definition** to open the Report Node Definition dialog box of the network node.

2 Click the **Sort** button to open the Sort dialog box.

3 Sort the column information, as follows:

Node	Sort item by	Then By	Order of Direction
network	network_netclass		Descending
		network_nettype	Descending
		Operation State	Descending
		network_domain	Descending
host	host_dsname		Descending
		operation state	Descending

Node	Sort item by	Then By	Order of Direction
service	Name	Name	Descending
		service_operatingstatus	Descending
		operation state	Descending

4 Click **OK** to close the Sort 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 Node Order

You must now determine the order in which the nodes appear in the report. Since the network node is the parent node, it should appear at the beginning of the report. The data of the service node is displayed next followed by the events related to the services.

To set the node order:

1 Right-click any node displayed in the Map pane 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 move the network node to the beginning of the list and the service node to the bottom.

The Set Node Order dialog box now displays as follows:
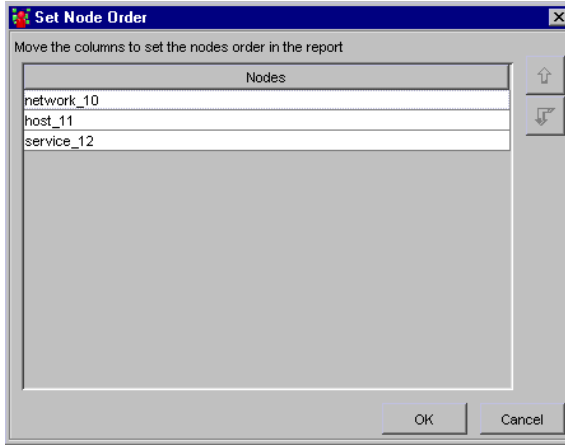
3 Click **OK**.

Refining and Saving the Report Layout

For better visual identification of the parent hosts and their dependent services, you define the host node as a child of the network node. The link between the service and host nodes is defined as sibling, to display their data side by side.

To refine and save the report layout:

1 In the Map pane, right-click the link in between the network and the host nodes and select **Child** to display the host node columns indented beneath the network node in the report.

2 Right-click the link between the service and the host and select **Sibling** to display the service node columns next to the host 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 On Report Manager, right-click the Services report and select **Generate**.

2 Select the HTML format to display and view the report in your browser.

3 Click **OK**. Click **OK** to the confirmation message.

4 Click **OK** to open the Report Download dialog box.

5 To display the report, select the **Open this report** button and click **OK**. The report is displayed, and is similar to the following figure:

6 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 241.

Part IX

The Class Browser

24

Introduction to the Class Browser

This chapter introduces the Class Browser that enables you to view the class model of your managed world.

This chapter describes:	On page:
Working with the Class Browser	369
Class Browser Shortcut Menus	367
Class Browser Links	368
Browsing the Class Model	369

About the Class Browser

The Class Browser is displayed by selecting the Class Browser tab in the Explorer pane. The Class Browser enables you to view the information in the class model, which consists of all the defined classes in the system and the links that define the relationships between them. The class model represents the structure of the managed world on which Mercury Application Mapping is used.

The class model visually reflects the CMDB, which contains the data on which the system, including TQL queries and managed views, is based.

Explorer pane   Maps pane
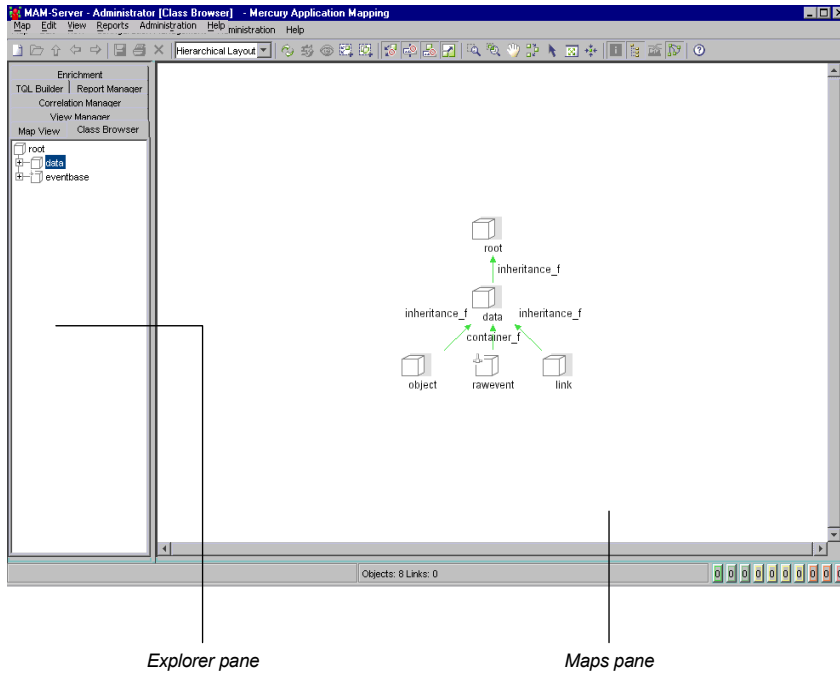
When the Class Browser tab is selected, the Mercury Application Mapping main window is divided as follows:

➤ **Explorer pane** – Displays a hierarchical tree structure of the class model containing the inheritance links between object classes. All classes included in the class model are classified in one of these two branches, either object or link. By selecting an object class in the Explorer pane, you can drill down and view the links and neighbors of the selected object in the Map pane.

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

➤ **Map pane –** Displays various layouts of the class model, including the classes and the link types connecting them. (For details on link types, see "Class Browser Links" on page 368.) Each class has a unique icon, its neighbors (as indicated by its connection to the class in the Explorer pane) all share the same shape surrounding the icon. (Shapes are predefined according to logical class groupings.)

Class Browser Shortcut Menus

The Class Browser contains a shortcut menu, which is displayed by right-clicking a selected class in the Map pane or in the Explorer pane. The shortcut menu provides quick access to the most commonly used functions in the Class Browser, as described in the following table:

Option	Description
Add/Remove Relationship	Enables you to add relationship links between classes, which define their physical/logical connections.
Edit Class	Enables you to view and edit the properties of the selected class, including its attributes, default tags, icon file, default labels, and so forth.
Show Class Instances	View inheritance links between classes – for details, see "Browsing the Class Model" on page 370
Export	Displays a Save As dialog box, enabling you to save the class as an XML script. This option can be used to move classes from one workstation to another provided the attached TQL query is also moved.
Hide	Enables you to hide the selected class.

Option	Description
Delete	Enables you to delete the selected class. Once you delete the class it no longer appears in the class model in the TQL Builder and Logical Object Manager. To restore it, you must recreate the database by using the **OnlinDBCreator** command file (for details, see the *Mercury Application Mapping Administration Guide*).
Get Neighbors	Enables you to display the classes that are connected to the selected class (that is, its neighbors) via a link.
Get Next Neighbors	(Appears only after using the Get Neighbors option.) Enables you to keep the previous Get Neighbors display on the Map pane, and to add the neighbors of one of the displayed objects to it.
Straight	(Appears only for links with angles on the Map pane.) Enables you to straighten links with angles.

Tip: Right-clicking an empty area in the Map pane displays a shortcut menu with the zoom and layout options available in the **View** menu, as described in "Mercury Application Mapping Quick Tour" on page 15.

Class Browser Links

The classes displayed in the Map pane of the Class Browser are connected by the following three link types:

➤ **Inheritance Links –** Functional links in which one class inherits the attributes of the class to which it is connected, while adding more attributes of its own. For example, an object in the host class might have its own attributes, such as DNS name and operating system, in addition to inherited attributes from the class above it, such as name and state.

Inheritance links are displayed both in the Explorer pane tree and as black links in the Map pane. As with all functional links, inheritance links have labels that end in **_f**.

➤ **Functional Links –** Functional links are used to describe a physical or logical connection between two classes which are not represented in the CMDB. (For details about functional links and their use, see "Functional Links" on page 62.)

Functional links are displayed in red (dashed lines) in the Map pane. All of them have labels that end in **_f**.

➤ **Relationship Links –** Links that describe a physical or logical connection between two classes. For example, an IP client is connected to an IP server via a client/server link. These links are displayed in green on the Map pane.

Working with the Class Browser

Use the Class Browser to perform the following functions:

➤ create new classes and edit existing ones

➤ expand or collapse the Explorer pane to display the required objects and links in the Map pane

➤ show or hide selected objects and links in the Map pane

➤ view the selected object's neighbors, which are connected using inheritance, container and other types of relationships.

➤ zoom in and out of the Map pane to view the selected object and its neighbors at different zoom levels

➤ display the class model using different layouts

Browsing the Class Model

The Class Browser enables you to browse the class model to better understand the relationships between specific classes.

This section includes the following topics:

➤ "Class Browser Tooltip" on page 370

➤ "Browsing the Class Model" on page 370

Class Browser Tooltip

When you point to an object in the Class Browser, the tooltip displays the following information:

Class – The name of the class to which the object belongs.

Instances – The number of objects that currently exist in the database from the selected class or its subclasses.

Browsing the Class Model

In the Explorer pane, you can view inheritance links between classes (such as the **Ipserver**, **Ipclient** and **Ipunknown** classes, which inherit attributes from the **IPport** class). In the Map pane you can view links of all types that connect various classes. If required, you can print the contents of the Map pane.

To browse the class model:

1 From the Explorer pane, select the class that you want to view in the Map pane. (The class is automatically selected in the Map pane.)
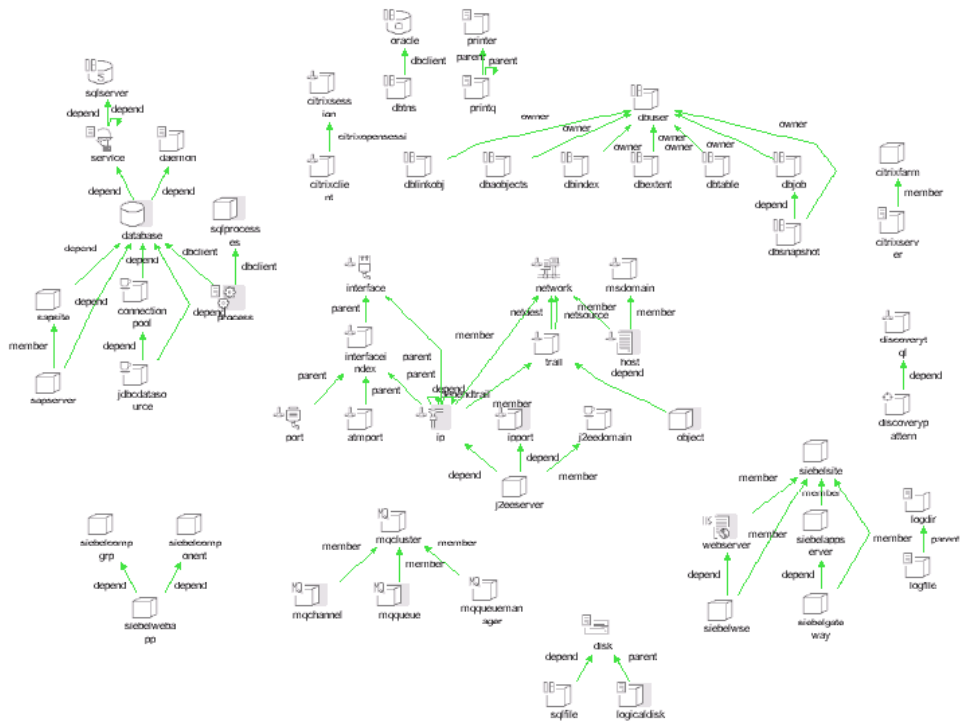
2 Use the interactive zoom button on the toolbar to zoom in on the view area of the selected layer.
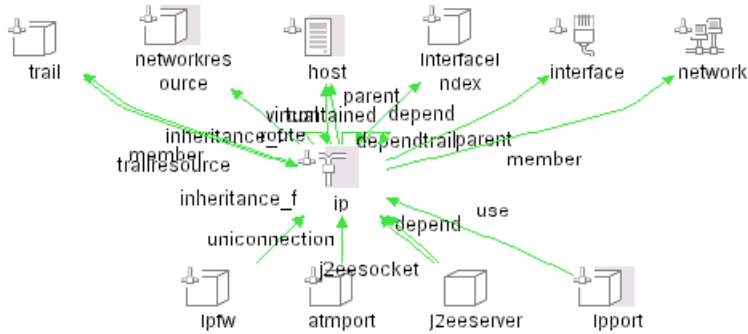
3 To return to the main view, click the **Fit To Window** button on the toolbar.

To view objects and links directly connected to the object:

1 Right-click an object and select **Show Class Instances.**

2 In the class instances table, right-click the object for which you want to see the objects and links directly connected to it and select **Get Neighbors** (for details, see "Displaying Interdependent Objects" on page 181).

25

Managing Classes

This chapter explains how to create and add classes to the class model.

About Managing Classes

If you have IT resources that are not included in the system's built-in classes, you can create and add classes to the class model. In addition, you can edit existing classes to adjust them to the structure of your IT infrastructure.

The workflow is as follows:

➤ Create classes (for details, see "Creating Classes" on page 374)

➤ Assign an icon to a class (for details, see "Assigning an Icon to a Class" on page 378)

➤ Create methods for the class (for details, see "Creating Class Methods" on page 380)

➤ Add tags to the class (for details, see "Adding Tags to the Class" on page 384)

➤ Define an attribute for the class label (for details, see "Defining an Attribute for the Class Label" on page 384).

Creating Classes

You can create and add classes to the class model.

To create a class:

1 Click the **Class Browser** tab.

2 Click **New** on the toolbar or select **New** in the **Map** menu. The Class Wizard is displayed.

3 In the **Class Name** box, enter a unique name for the new class. Do not use a blank space or an underscore as part of the class name. You can use lower and upper case, but you cannot use the same name with different cases for two classes.

4 In the **Base Class** box, select a base class for the class you are creating. The new class inherits the base class' attributes.

5 (Optional) In the **Class Description** box, enter a description of the new class.

6 (For links only) If you do not want the relationship between the two objects to have a specific direction, select **Bi-directional Link**. A TCP link between two servers, for example, can be defined this way.

7 Click **Next**. The Class Attributes dialog box is displayed.

The **Class Attribute** box displays all the attributes that the new class inherits from its base class.

8 To add attributes, click **Add** to open the Add Attribute dialog box.

9 Enter information about the new attribute according to the following information:

➤ **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 of the new attribute.

➤ In the **Attribute Type** area, select one of the following:

- **Primitive –** Choose from one of the following field types: **Boolean**, **Date**, **Double**, **Integer**, **Long**, **Password**, or **String**.

- **Class –** Contains a list of all classes.

- **Enumeration/List –** Contains a list of Lists/Enumerations defined in the Enumeration Manager dialog box. 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

For details on the Enumeration Manager, see Chapter 26, "Enumerations and Lists."

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 class is created and there is no runtime value for the attribute.

➤ 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 class.

- (Optional) Select **Visible** to display this attribute in the Object Attributes dialog box in the Map View.

- (Optional) Select **Change Monitored** to define attributes whose values are being marked as **Change Monitored** in the Class Browser. For every change in an attribute defined as **Change Monitored**, you are notified by a Change event in the Map View. For details on how to track modifications made in the Map View, see "Viewing and Editing Object Attributes" on page 195.

- (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 Attribute dialog box.

- (Optional) Select **Save History** to keep a record of the attribute history in the Mercury Universal CMDB.

- (Optional) Select **Comparable** to enable this attribute to be used for comparing compound objects. For information on how to compare compound objects, see "Comparison Capabilities" on page 411.

- (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 264.

 10 Click **OK** to save the changes you have made. The new class attribute you created appears in the Class Attributes dialog box.

Note: If you modify an attribute belonging to its parent, it turns light blue. If you modify an attribute belonging to the class itself, it turns dark blue.

---

link_end1	End1	object		
link_end2	End2	object		

11 If you want to add more attributes, repeat steps 8 to 10.

12 (Optional) 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:
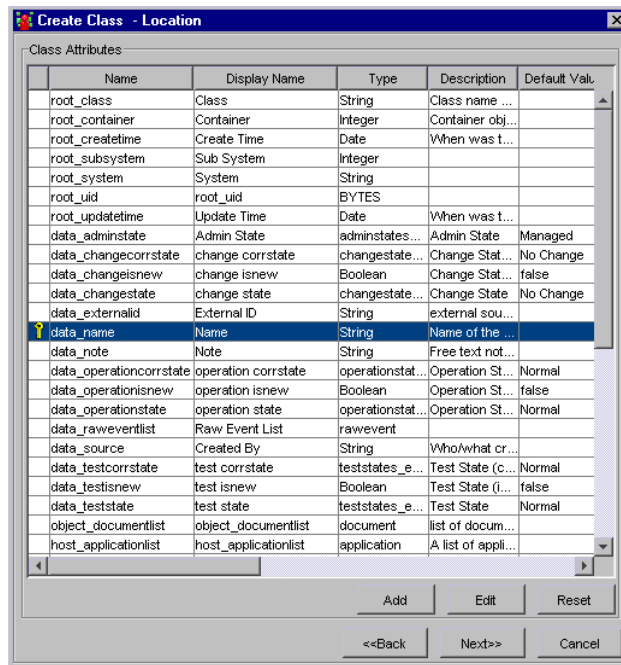
13 Click **Next** to attach an icon to a new class (for details, see below"Assigning an Icon to a Class" on page 378).

Assigning an Icon to a Class

Mercury Application Mapping enables you to attach an icon to a new class. You can also attach different icons to the same class when certain conditions apply. For example, you can associate different icons with the same class when one of its attribute values changes.

Note: You can only assign different icons to classes that are classified as objects, and not links.

To assign one or more icons to a class:

1 In the Class Attributes dialog box (for details, see step 7 in "Creating Classes" on page 374), click **Next**.

2 From the **Class Family** list, select the icon you want to attach to the specified class (for details, see "Object Level" on page 42).

3 From the **Class Main Icon** list, select the group the class belongs to.

4 Select **Advanced** to assign different icons when the selected class 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 new 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%**.
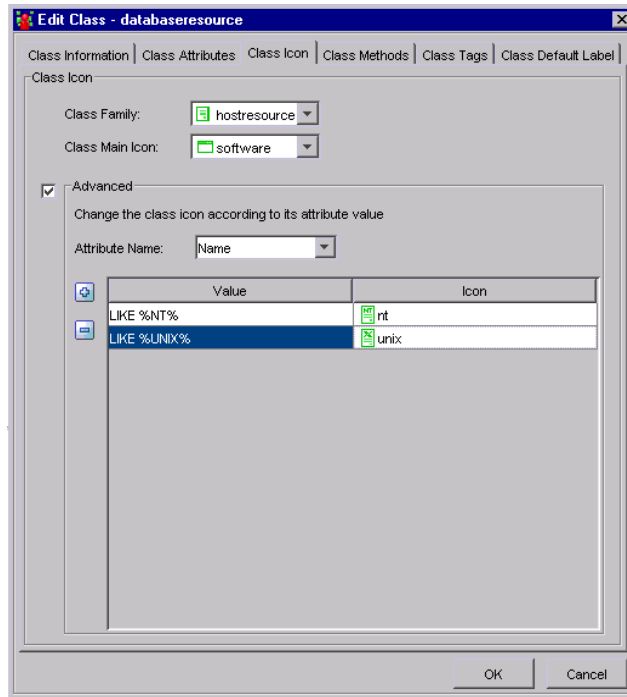
Note: To delete a row, select the row you want to delete and click the minus sign.

10 Click **Next** to create new class methods.

Creating Class Methods

Mercury Application Mapping enables you to customize the methods (commands) and shortcut menu of a class in the Class Browser. 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 object in the Map View.

To add a menu item:

1 In the Class Attribute dialog box (for details, see "Assigning an Icon to a Class" on page 378), click **Next** to display the Class Methods dialog box.
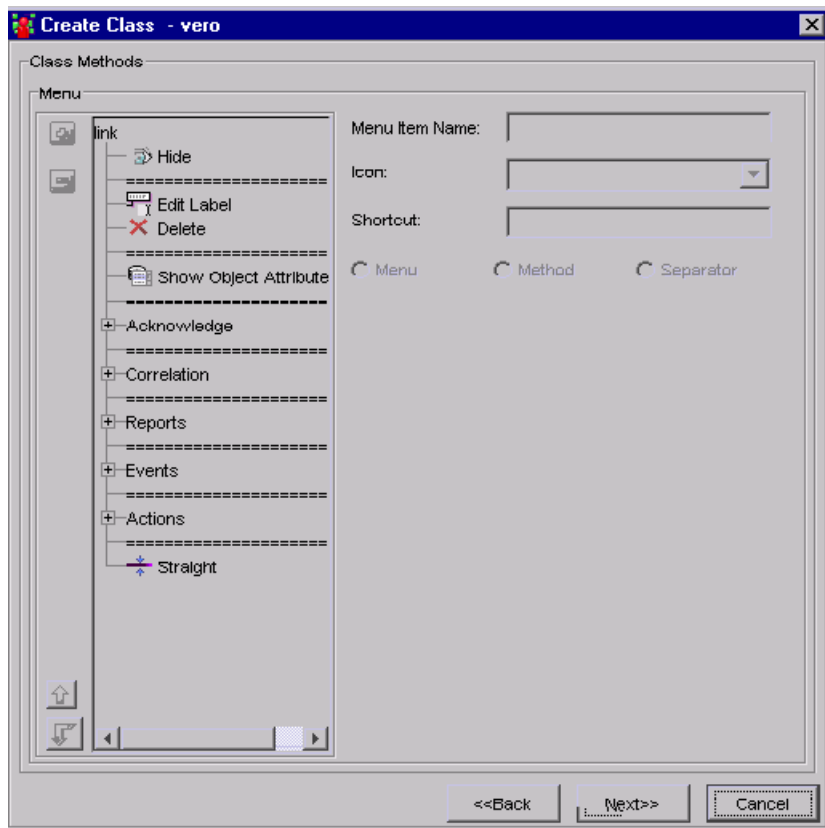
The field on the left displays a hierarchical tree containing the default menu items for the selected class.

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 Map View.

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 Root class and all its ancestors.

> **Tip:** To create a method and not override the existing method in the
> selected class, 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 object 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 object attributes in the **Parameters** box, do the following:

➤ Click the plus button to create a new entry.

➤ Type the object's attribute.

> **Note:** To delete an existing entry, select it and click the minus sign.

8 To place a separator between two menu items, do the following:

➤ Click **Separator**.

➤ Click **OK** to place a separator underneath the selected menu option as shown in the illustration below.

separator
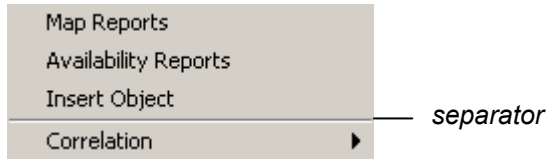
9 (Optional) In the **Confirm Message** box, type the message you want to appear when you select the menu item. Select the check box to the left of the **Confirm Message** box to activate this option.

Note: To delete a menu item, select an item and click the minus sign.

---

10 To change the order of the menu items, drag and drop an item into the required position.

Note: If a class does not have a method defined specifically for it, the class inherits all the menus from its parent class 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 class that is edited.

---

11 Click **Next** to add tags to the class (for details, see "Adding Tags to the Class" on page 384).

Adding Tags to the Class

This section describes how to add tags to the new class.

To add tags to a class:

1 In the Class Methods dialog box (for details, see "Creating Class Methods" on page 380) click **Next** to open the Class Tags dialog box.

2 From the **Available Tags** box, select the tags that you want to add to the new class and click the left-to-right arrow. The selected tags are displayed in the **Class Tags** box.

3 (Optional) If required, you can define new tags by entering a name for the tag in the space provided and clicking **Insert**. The new tag that you entered is displayed in the **Class Tags** box.

4 To add all available tags to the selected object's tags, click the double left-to-right arrow.

5 To delete a tag from the Selected Object Tags area, select it and click **Delete**.

Note: You cannot edit existing tags.

6 Click **Next** to add an attribute to the class label (for details, see "Defining an Attribute for the Class Label" on page 384).

Defining an Attribute for the Class Label

You can define an attribute for the class label.

To define an attribute for the class label:

1 In the Class Tags dialog box (for details, see "Adding Tags to the Class" on page 384), click **Next**. The Class Default Label page is displayed (for details on how to define an attribute for the class label, see "Formatting Text" on page 101).

Note: If you defined a Label Format for the class in the View Manager, it overrides the default labels you define here. For details, see "Adding Attributes to the Label of the Selected Object" on page 101.

---

2 Click **Finish**. Mercury Application Mapping creates the new class using the information you provided, and displays it on the Explorer pane in the class model.

Editing Existing Classes

You can edit an existing class.

To edit an existing class:

1 In the Explorer pane, right-click the class you want to edit and click **Edit Class** to open the Edit Class dialog box.

The Edit Class dialog box tabs are similar to the Class Wizard (as described in "Creating Classes" on page 374). However, only editable fields are enabled in the Edit Class dialog box.

2 Make your changes in the different tabs as follows:

- ➤ Class Information (for details, see "Creating Classes" on page 374).
- ➤ Class Attributes (for details, see "Creating Classes" on page 374).
- ➤ Class Icon (for details, see "Assigning an Icon to a Class" on page 378).
- ➤ Class Methods (for details, see "Creating Class Methods" on page 380).
- ➤ Class Tags (for details, see "Adding Tags to the Class" on page 384).
- ➤ Class Default Label (for details, see "Defining an Attribute for the Class Label" on page 384).

Importing Classes

You can import XML files that contain saved classes to your Class Browser. Use this option if you want to move classes from one workstation to another.

To import a class:

1 From the Class Browser, open the **Map** menu and click **Import** to open the Import dialog box (for details, see "Importing a TQL Query" on page 74.

2 Locate the class you want to import, select it and click **Import**. The imported class is added to the class model in the Explorer pane.

Adding Relationship Links Between Classes

You can add relationship links between classes, either built-in or new ones, which define their physical or logical connections.

To add a relationship link between two classes:

1 In the Map pane, select the two classes you want to connect by holding down CTRL and clicking the class names.

2 Right-click one of the classes and select **Add/Remove Relationship** to open the Add Relationship dialog box.

3 Select the check box(es) that specifies the type of link you want to create.

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

26

Enumerations and Lists

This chapter explains how you create a predefined list whose values define an attribute type.

This chapter describes:	On page:
About Enumerations and Lists	387
Creating an Enumeration or List Definition	388
Deleting List and Enumeration Type Definitions	389

About Enumerations and Lists

Mercury Application Mapping enables you to create a predefined list whose values define an attribute type. You can create a definition for the following attribute types:

➤ **List** – Enables you to create a predefined list of values. For example, Location.

➤ **Enumeration** – Similar to List but has more capabilities, such as assigning a color for every value. Enumerations are designed to be used as severity value lists by categories. For details on categories, see the *Mercury Application Mapping Administration Guide*.

Creating an Enumeration or List Definition

You can create an enumeration or list definition.

To create a list definition:

1 From the **Administration** menu, select **System Type Manager**.

2 Click the **Add** button.

3 In the **Name** box, type a unique name for the new type definition.

4 To create a definition for the type List, select the **List** option.

5 From the **Type** box, choose a field type: **Date**, **Double**, **Integer**, **Long**, or **String**.

6 Click the plus button to add a row.

7 Double-click inside the row and either select a value from the list that appears or type the required value.

8 To add another value, repeat steps 6 and 7.

9 Click **OK**. The new List type definition you have created appears in the Enumeration Manager dialog box.

To create an enumeration definition:

1 From the **Administration** menu, select **System Type Manager**.

2 Click the **Add** button to open the Create Enumeration Definition dialog box.

3 In the **Name** box, type a unique name for the new type definition.

4 To create a new Enumeration, select the **Enumeration** option.

5 Click the plus button to create a new row.

6 In the **Value** box, type a value. It can be a string or a number; for example, **Red** or **my value**. The value appears in the tooltip for the object in the Map View.

7 In the **Key** box, type a number to create an Enumeration that describes a severity list for a category. For example:

Key	Value	Severity it represents
0	Green	Normal
7	Orange	Major
10	Red	Critical

Assign key values according to the following rules:

➤ The list of key values must always begin with zero (0). (Zero is used to represent a Normal state.)

➤ The list must always be numbered consecutively.

8 Select a color that indicates the severity level in the **Color** area.

9 To add another entry in the severity list, repeat steps 5 to 8.

10 Click **OK**. The attribute type you have defined is displayed in the Enumeration Manager dialog box.

Deleting List and Enumeration Type Definitions

You can delete List or Enumeration type definitions.

To delete List and Enumeration definitions:

1 From the **Administration** menu, select **System Type Manager**.

2 Click the **Add** button to open the Create Enumeration Definition dialog box.

3 To delete a row for both List and Enumeration type definitions, select the row you want to delete and click the collapse button.

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

Part X

Mercury Application Mapping Applications

27

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.

This chapter describes:	On page:
Comparing Snapshots Taken at Different Times	396
Understanding the Compare Snapshots Dialog Box	399

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 394

➤ "Saving Snapshots on Periodic Basis" on page 394

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 Explorer pane of the Map View, right-click the view of which you want to take a snapshot.

2 Select **Snapshot** > **Save** to open the Save Snapshot dialog box.
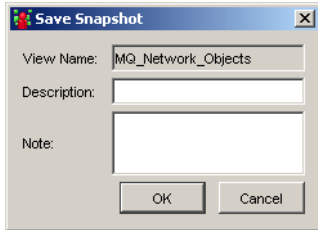
3 The **View Name** box contains the name of the view as it appears in the Explorer pane. You cannot edit the view's name.

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 in the Mercury Application Mapping status bar:

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.
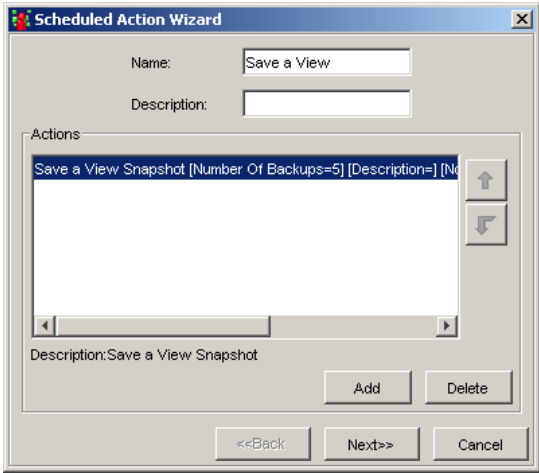
For information on how to schedule the task, see the *Mercury Application Mapping Administration Guide*.

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.
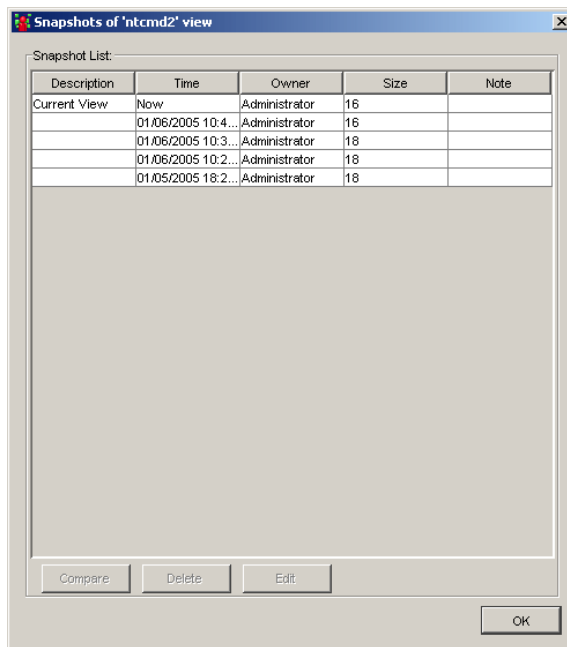
This dialog box contains a list of all the different snapshots taken of the selected view.

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 394).
Time	The time at which the snapshot was saved.
Owner	The name of the user who took the snapshot.
Size	The number of objects 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 394).

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.
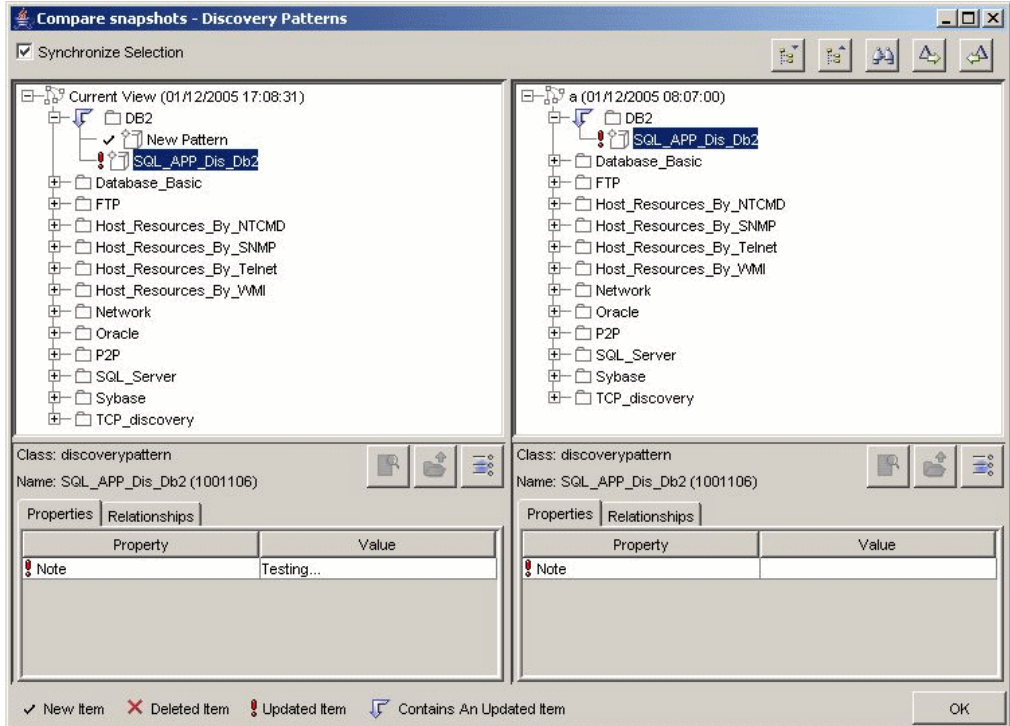
Note: For details about the Compare snapshots dialog box, see "Understanding the Compare Snapshots Dialog Box" on page 399.

To edit a snapshot:

 1 Right-click the view for which you have taken snapshots.

 2 Select **Snapshot** > **View snapshots** to open the Snapshots dialog box.

 3 Select the snapshot you want to edit from the **Snapshot List** section on page 396.

 4 Click the **Edit** button to open the Edit existing snapshot dialog box

 5 Edit the dialog box according to the procedure on page 394.

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

To delete a snapshot:

1 Right-click the view for which you have taken snapshots.

2 Select **Snapshot** > **View snapshots** to open the Snapshots dialog box.

3 Select the snapshot you want to delete from the **Snapshot List** and click the **Delete** button.

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 objects in the map view for that snapshot version. The left side always represents the newer of the two versions. When you select an object on one side, the corresponding object 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 an object:

Class – The name of the class to which the object belongs as defined in the Class Model.

Name – The name of the object as defined in the object's label.

The rest of this section contains the following topics:

➤ "Compare Snapshots Dialog Box Tabs" on page 400

➤ "Expanding and Collapsing the Entire Tree" on page 400

➤ "Finding Items Marked as Changed" on page 401

➤ "Finding an Object by Its Label" on page 401

➤ "Change Indicators" on page 402

➤ "Viewing and Comparing a Configuration File" on page 405

➤ "Displaying an Object's Attributes" on page 408

➤ "Enabling Synchronized Selection" on page 409

Compare Snapshots Dialog Box Tabs

The Compare snapshots dialog box contains the following two tabs:

Tab	Description
Properties	Displays only the attributes of the object whose values have been defined as **Change Monitored**. For more information about the **Change Monitored** option, see "Receiving Change Events When a Change Monitored Attribute is Modified" on page 233. The Properties tab has the following two fields: ➤ **Property** – Displays the attribute name for the selected object which was marked as **Change Monitored** in the Class Browser. ➤ **Value** – Displays the value for the currently selected object.
Relationships	Displays the links with which the different objects are connected within the view. The Relationships tab has the following two fields: ➤ **Property** – Displays the link 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 object 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 an Object by Its Label

Mercury Application Mapping enables you to search for an object within the selected view whose label contains the searched text. You can specify a whole word or part of it.

To find the object by its label:

1 Click the **Find** button to open the Find object by label dialog box.
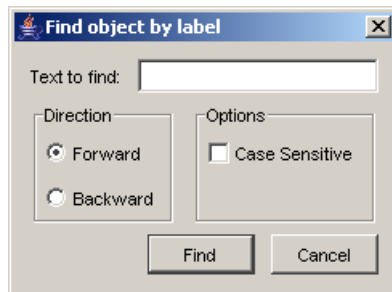
2 In the **Text to find** box, type the text you want to find.

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 objects that appear in the map view 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 object has been added to the view. The example below shows that the object **aPPaPPP** has been added to the view in the snapshot at the left.
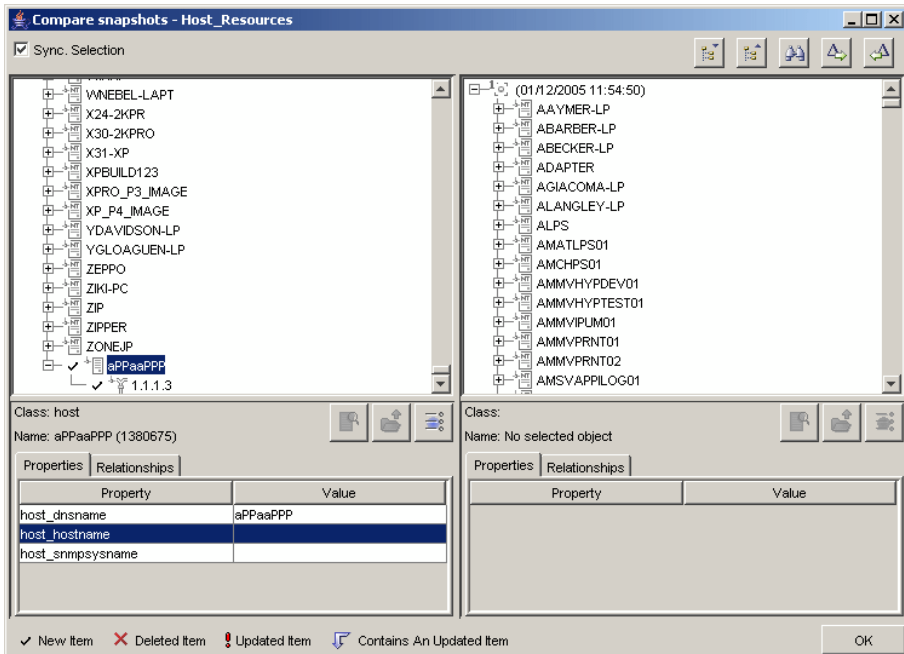
Deleted Item Icon

The **Deleted Item** icon indicates that an object has been deleted from the view. For example, the figure below shows that the object **DR2** has been deleted from the view in the snapshot at the right.
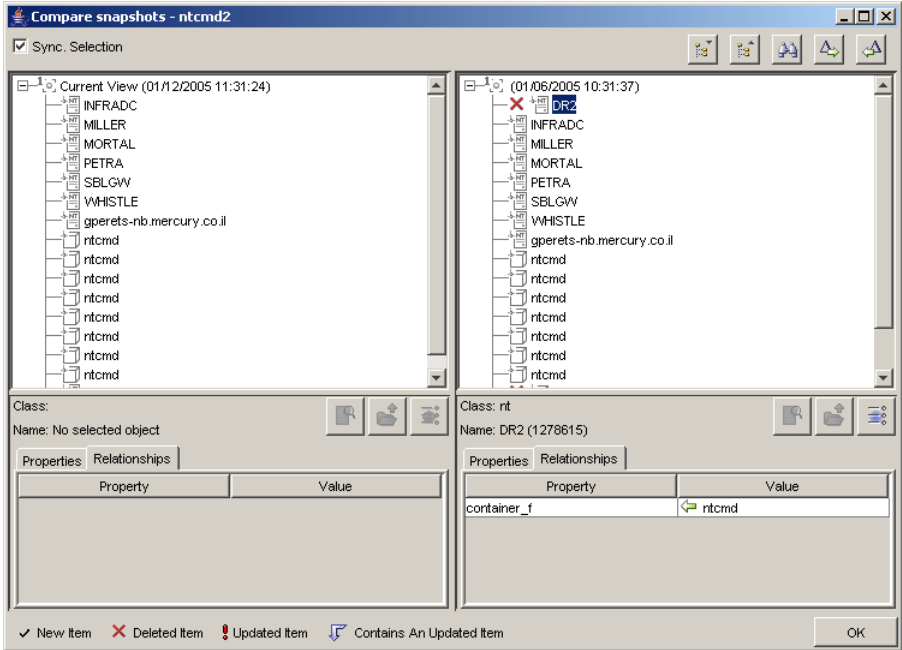
Updated Item Icon

The **Updated Item** icon indicates that an object whose attribute value has been defined as **Change Monitored**, has changed For example, the figure below shows that the object 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).
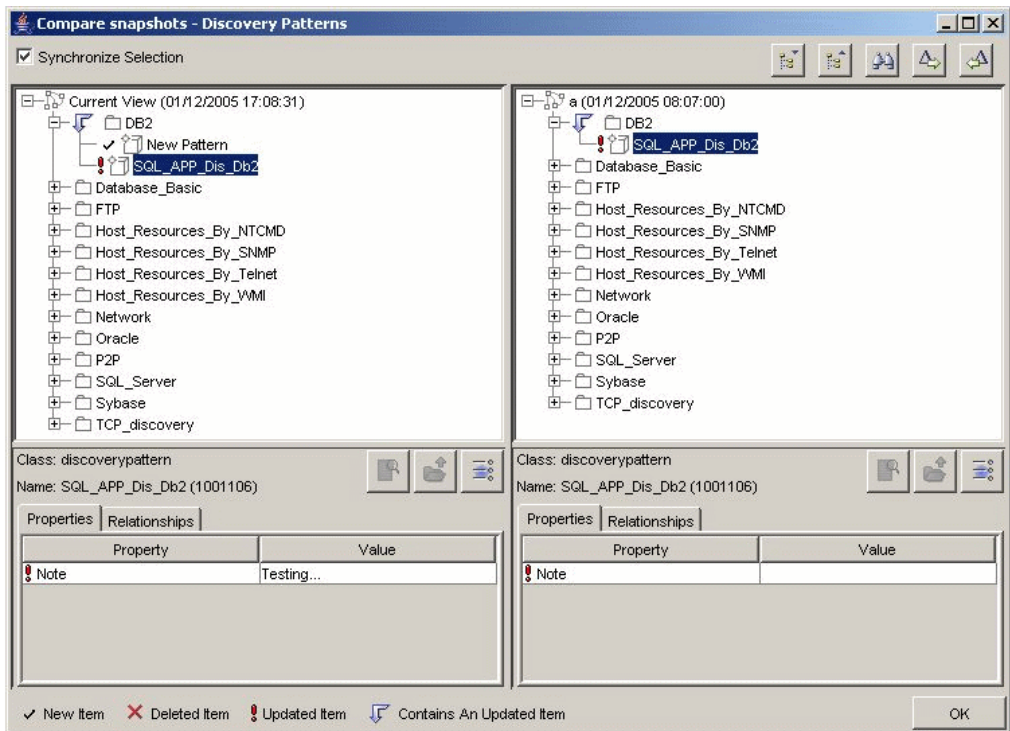
Contains an Updated Item Icon

The **Contains an Updated Item** icon indicates that either a new child object was created for an object or an existing child object was edited or deleted. For example, the **DB2** object in the figure above has a **Contains an Updated Item** icon due to the changes made to its child objects.

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 an object of the class **configfile** from the tree.
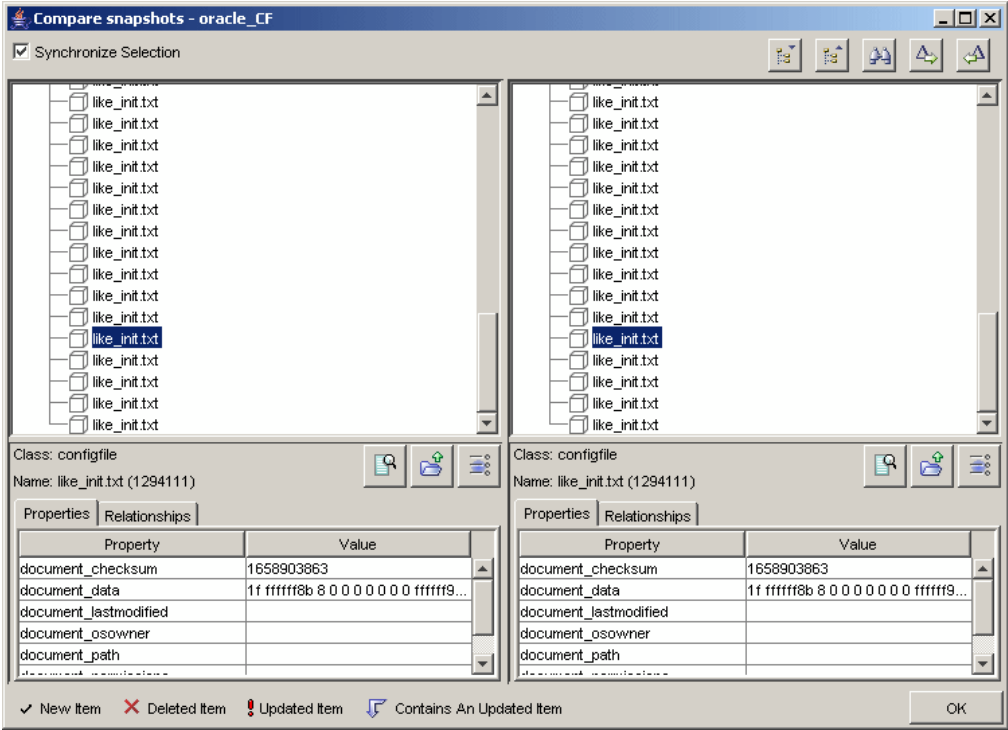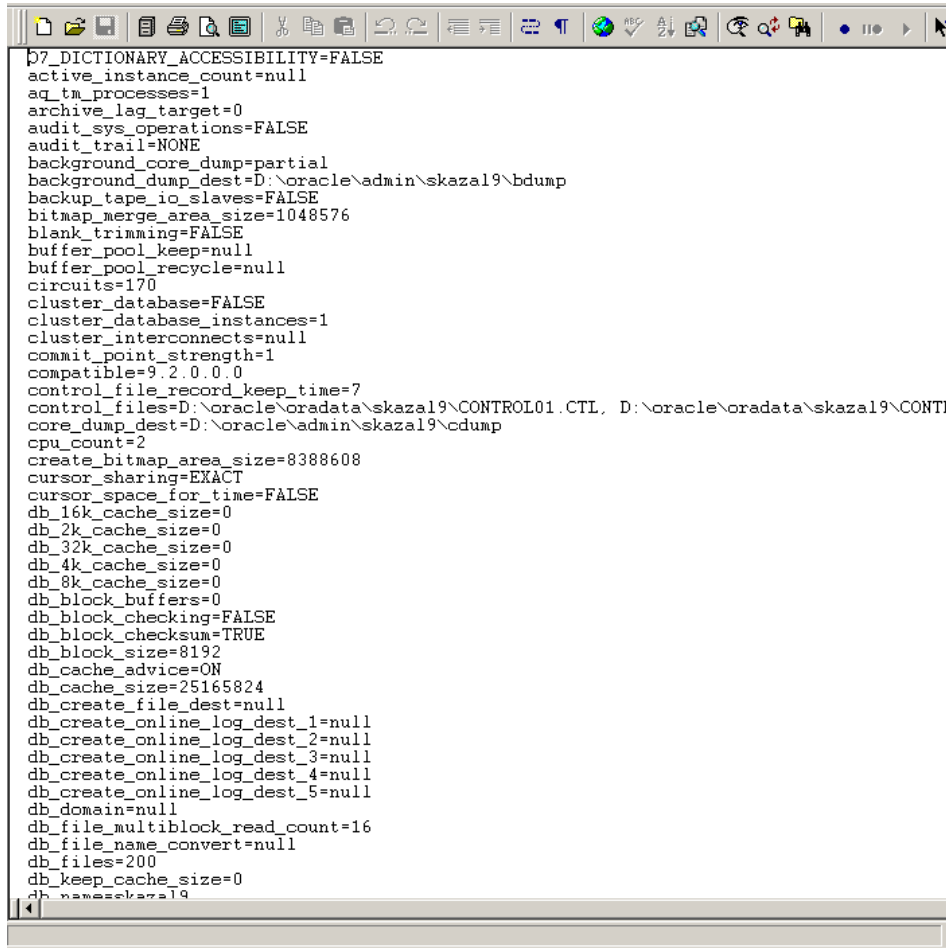
 2 Click the **Open** button to open the configuration file.

Note: This button is only enabled if you have selected an object of the class **configfile** from the tree.

To compare what the configuration file looked like at different times:

1 In the Compare Snapshots dialog box, select an object of the class **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 an object of the class **configfile** from the tree.

The differences in the configuration file are marked in blue.

Displaying an Object's Attributes

You can display what the object's attributes were at the time the snapshot was taken.

To display an object's attributes:

 1 In the Compare Snapshots dialog box, select an object from one of the trees.

 2 Select the **Show Object Attributes** button to open the Object Attributes dialog box.
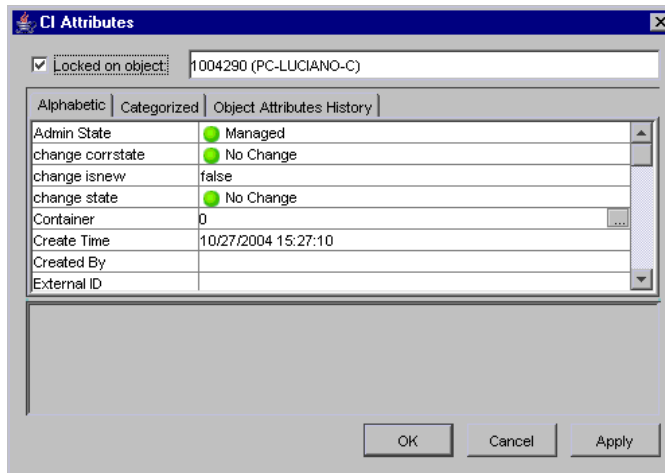
Note: This dialog box is for viewing only.

The Object Attributes dialog box lists the classes contained in this view and their attribute values.

 ➤ The classes in the **Alphabetic** tab are listed in alphabetical order.

 ➤ The classes in the **Categorized** tab are grouped according to the categories defined in the Class Browser.

 3 If you are using the **Categorized** tab, click the **Plus** button to view all the classes contained in this view.

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 an object on one side and have its corresponding object automatically selected on the other side.

By default, **Synchronize Selection** is selected. Clearing this check box allows you to compare two different objects by selecting one from each of the two trees.

28

Comparison Capabilities

This chapter describes Mercury Application Mapping's comparison capabilities.

This chapter describes:	On page:
Gold Master Comparison	411
Object Comparison	420

Gold Master Comparison

Mercury Application Mapping enables you to generate a report that compares the configuration of a Gold Master object to other objects of the same class.

You choose the object you want to use as the Gold Master object. Then you choose the objects whose configuration you want to compare to the Gold Master template. You can choose the objects from existing views or from the CMDB.

In the report that is generated, Mercury Application Mapping specifies which of the selected objects are configured differently.

Note: The objects in this report will only appear if their attributes are marked as **Comparable** in the Edit Attribute dialog box (see "Creating Classes" on page 374.)

This section includes the following topics:

➤ "Comparing the Configuration of Objects to a Gold Master Template" on page 412

➤ "Adding Multiple Objects to the Object Items List" on page 416

➤ "Deleting a Gold Master Report" on page 417

➤ "Editing a Gold Master Report" on page 417

➤ "Displaying an Object's Attributes" on page 417

➤ "Searching for an Object" on page 418

➤ "Displaying Related Objects" on page 419

Comparing the Configuration of Objects to a Gold Master Template

This section describes how to compare the configuration of specific objects to a Gold Master template.

To compare the configuration of objects to a Gold Master template:

1 In Map 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.
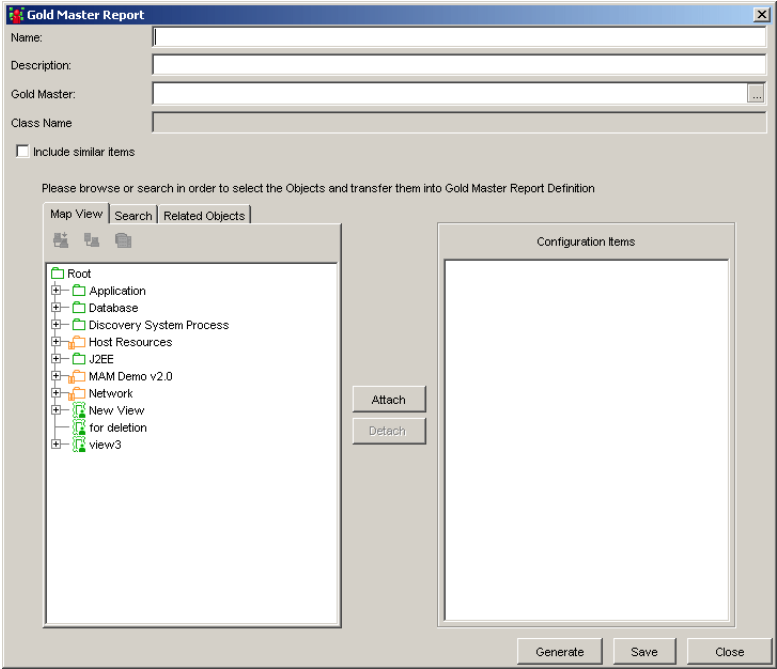
Note: You can also open the Gold Master Report window by right-clicking an object and selecting **Gold Master Report**.

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 object you want to use as a Gold Master object, click the button at the right end of the **Gold Master** box.

6 In the Map pane or the View Explorer, select the required object. The selected object appears in the Chosen object box, as seen below:

7 Click the **Choose Object** button. The Gold Master Report window opens and the object you selected as your Gold Master template appears in the **Gold Master** box.

The Gold Master object's class appears in the **Class Name** box.

8 Select the **Include similar items** check box if you want to include the compared objects whose configuration is identical to that of the Gold Master object.

9 In the left pane, select the object(s) that you want to compare to the Gold Master object. You can select the objects in the following ways:

➤ In the **Map View** tab, select the required object(s) at different levels of the view hierarchy. You can select multiple objects by holding down the CTRL button and clicking the required objects. To view and edit an object's attributes, see "Displaying an Object's Attributes" on page 417.

➤ In the **Search** tab, search for objects in existing views or in the entire CMDB. For details, see "Searching for an Object" on page 418.

➤ In the **Related Objects** tab, select an object that is related to the selected object in the **Map View** tab. For details, see "Displaying Related Objects" on page 419.

10 Use the **Attach** and **Detach** buttons to transfer to and remove objects from the **Object Items** list. You can also select multiple objects. The objects that appear in the **Object Items** list are the objects that are compared to the Gold Master object.

11 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.

12 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**

13 Click **OK**. The Report Download dialog box opens.

➤ To open the report, select **Open This Report**.

➤ To save this report to disk, select **Save this report on disk** and browse to the desired location in the Save dialog box.

➤ 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 Object	The object that is used as the template to which to compare the other objects.
Report Time	The time the report was created.
Component	The class to which the object belongs.
Parent Object	The object under which the compared object is located.
Count/Gold	Number of objects in the compared objects as opposed to the number of objects in the Gold Master object.
Missing	Lists the objects that exist in the Gold Master object but not in the compared object.
Additional	Lists the objects that exist in the compared object but not in the Gold Master object.

Adding Multiple Objects to the Object Items List

You can add multiple objects to the **Object Items** list by first selecting the required objects in Map pane or the View Explorer and then opening the Gold Master Report window.

To add multiple objects to the Object Items list:

1 In Map pane or the View Explorer, select the required objects. You can select multiple objects by holding down the CTRL button and clicking the required objects.

2 Right-click and select **Gold Master Report** to open the Gold Master Report dialog box. The selected objects appear on the **Object items** list.

3 Use the to **Attach** and **Detach** buttons to add and remove objects from the **Object Items** list.

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 Objects to a Gold Master Template" on page 412.

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

Displaying an Object's Attributes

This section describes how to display and edit an object's attributes.

To display and edit an object's attribute:

1 In the Gold Master Report window, select the **Map View** tab.

2 Select the object whose attributes you want to view.

3 Click the **Show object's attributes** button to open the Object Attributes dialog box. The Object Attributes dialog box enables you to view and edit the attributes defined for the selected object. For details, see "Viewing and Editing Object Attributes" on page 176.

Searching for an Object

Mercury Application Mapping search capabilities enable you to search for objects to compare to the Gold Master object either in specific views or in the entire CMDB. Theses capabilities also provide different search criteria, through which you can search for objects according to their class, label, and/or attributes.

To search for objects:

1 In the Gold Master Report window, select the **Search** tab.

2 From the **Class** list, select the class of the object(s) you want to find.

3 (Optional) Select **Derived** to find both the selected class object(s) and any objects derived from the selected class (inheritance objects).

4 (Optional) In the **Options** area, if you want to find a specific object(s) whose label you know, enter the object'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.

5 (Optional) In the **Advanced** area, to search for an object(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:

• **Name** – Select the required attribute from the list.

• **Operator** – Select the required operator (such as =) from the list (for details, see "Operator List" on page 45)

• **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** – Select **And** or **Or** to link multiple conditions.

➤ After you complete the attribute definition, you can change it by using the **Edit** and **Delete** buttons.

6 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.

7 Click **Search**. The search results are displayed in the **Search** tab.

8 To do another search, click **Back to Search**.

Displaying Related Objects

Mercury Application Mapping enables you to display the interdependencies of any selected object (for details, see "Displaying Interdependent Objects" on page 181). You can select objects to compare to the Gold Master object from a list of related objects.

You can display:

➤ All objects connected to a selected object that are part of a specific view.

➤ All objects connected to the selected object in the entire CMDB.

You select an object in the **Map View** tab, and the objects related to the selected object appears in the **Related Objects** tab.

To display the related objects of a selected object:

1 In the Gold Master Report window, select the **Map View** tab.

2 Select the object whose related objects you want to display.

3 To display related objects, do the following:

➤ Click the **Get Related Object from View** button to display related objects in the selected view.

➤ Click the **Get Related Object from Database** button to display related objects in the CMDB.

The related objects are displayed in the **Related Objects** tab.

Object Comparison

Mercury Application Mapping enables you to:

➤ Compare compound objects. For details, see "Comparing Compound Objects" on page 421.

➤ Compare objects in a selected view. For details, see "Comparing Objects by View" on page 424.

➤ Understand the Compare Objects Dialog box. For details, see "Understanding the Compare Objects Dialog Box" on page 425.

➤ Understand the Properties tab. For details, see "Understanding the Properties Tab" on page 425.

➤ Expand and collapse the Tree. For details, see "Expanding/Collapsing the Tree" on page 425.

➤ Find items marked as changed. For details, see "Finding Items Marked as Changed" on page 426.

➤ Find objects by their label. For details, see "Finding an Object by its Label" on page 426.

➤ Understand change indicators. For details, see "Understanding Change Indicators" on page 427.

➤ View and compare configuration files. For details, see "Viewing and Comparing Configuration Files" on page 428.

➤ Display an object's attributes. For details, see "Displaying an Object's Attributes" on page 432.

➤ Enable synchronized selection. For details, see "Enabling Synchronized Selection" on page 433.

Comparing Compound Objects

A compound object is an object that is linked to another object by a **container_f** link. Comparing two compound objects enables you to see the difference between the hierarchies of the two objects. 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 Class Browser (see "Creating Classes" on page 374). For example, there might be a difference in the operating system or disk capacity.

Note: You can only compare two objects from the same class.

When you compare two compound objects, you can either first select the objects from the Map pane and then open the Compare compound objects window, or, alternatively, open the Compare compound objects window first and then select the two objects you want to compare.

Mercury Application Mapping also allows you to display the child objects of the compared objects that appear in the selected view. For details, see "Comparing Objects by View" on page 424.

To compare compound objects:

1 To compare compound objects, do one of the following:

➤ In Map View or View Explorer, select the two objects you want to compare. Right-click and select **Actions** > **Compare Compound** to open the Compare compound objects from DB dialog box:
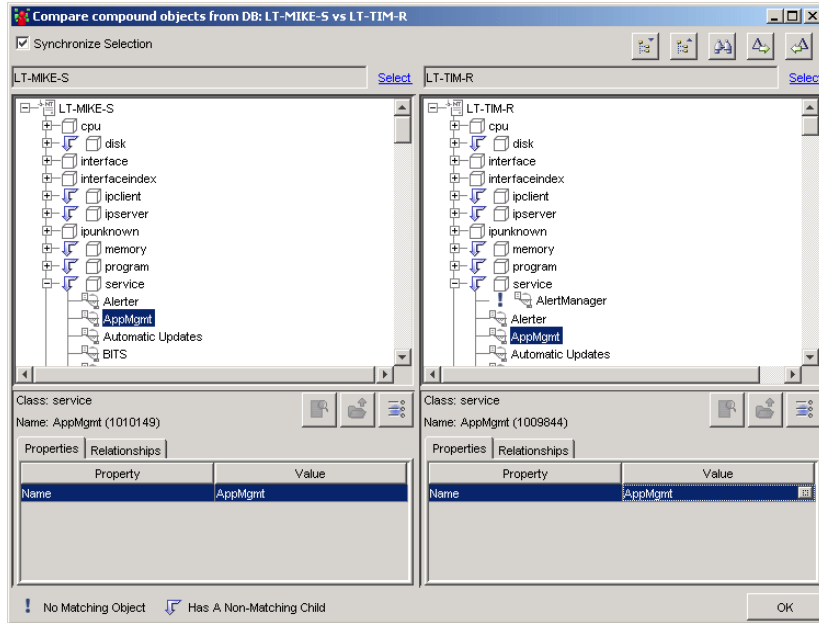
➤ Select **Configuration Management** > **Compare Compound** to open an empty Compare compound objects from DB dialog box.

Select

2 Click **Select** above the top-right corner of each pane to select the objects you want to compare. The following is displayed.
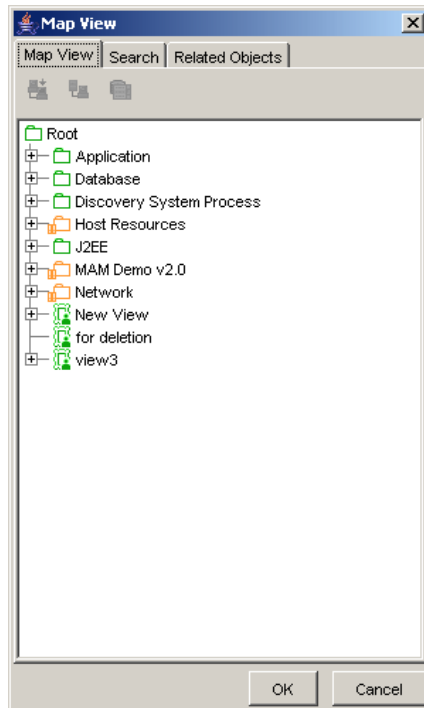
You can select the objects in the following ways:

➤ In the **Map View** tab, select the required object(s) at different levels of the view hierarchy. You can select multiple objects by holding down the CTRL button and clicking the required objects. To view and edit an object's attributes, see "Displaying an Object's Attributes" on page 417.

➤ In the **Search** tab, search for objects in existing views or in the entire CMDB. For details, see "Searching for an Object" on page 418.

➤ In the **Related Objects** tab, select an object that is related to the selected object in the **Map View** tab. For details, see "Displaying Related Objects" on page 419.

Comparing Objects by View

Comparing objects by view allows you to display the compared objects and their child objects that appear in the selected view.

To compare objects by view:

1 In Map View or View Explorer, select the two objects you want to compare.

2 Right-click and select **Actions** > **Compare By View** to display the selected objects and their child objects that appear in the view. The Compare objects by view dialog box opens.
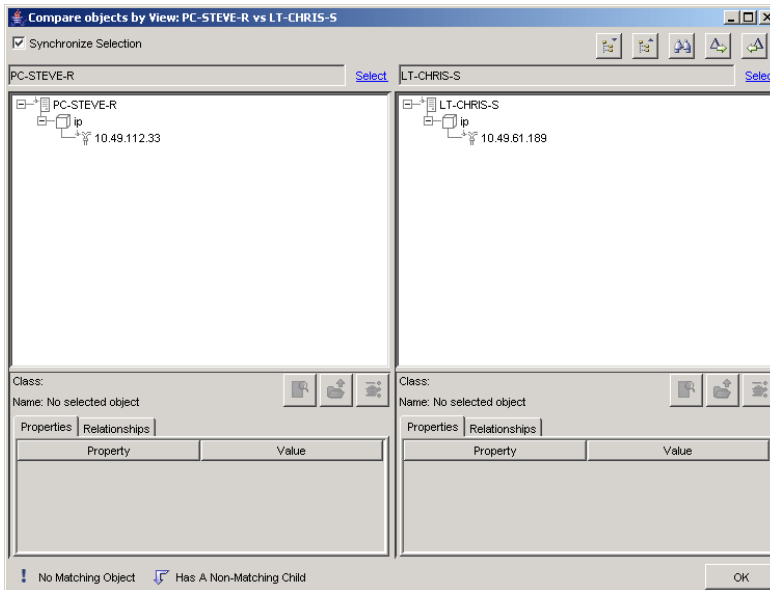
Understanding the Compare Objects Dialog Box

The Compare Objects dialog box is divided into two sections. Each side displays an hierarchical tree structure of the objects you selected. When you select an object on one side, the corresponding object on the other side is automatically selected, provided the **Synchronize Selection** check box is selected. (For details, see "Enabling Synchronized Selection" on page 433.)

The following information is displayed under the tree when you select an object:

Class – The name of the class to which the object belongs as defined in the Class Model.

Name – The name of the object as defined in the object's label.

Understanding the Properties Tab

The **Properties** tab displays only the attributes of the object whose values have been defined as **Comparable**. For more information about the **Comparable** option, (see "Creating Classes" on page 374.

The **Properties** tab has the following two fields:

Field	Description
Property	Displays the attribute name for the selected object which was marked as **Comparable** in the Class Browser.
Value	Displays the property value for the currently selected object.

Expanding/Collapsing the Tree

This section describes how to expand and collapse the tree.

To expand the entire tree:

Select an object and click the **Expand All** button.

To collapse the tree:

Select an object 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 427.

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 an Object by its Label

You can search for an object within the selected view whose label contains the searched text. You can specify a whole word or part of it.

To find the object for which you are looking:

1 Click the **Find** button to open the Find object by label dialog box.
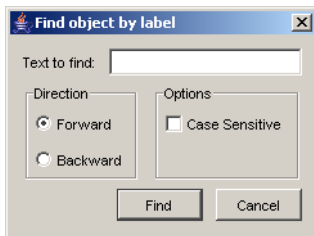
2 In the **Text to find** box, type the text you want to find.

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 objects.

Understanding Change Indicators

This section describes the icons that indicate the kind of change that has occurred in the compared objects. The indicator key appears at the bottom of the Compare compound objects dialog box.

This section includes:

➤ "No Matching Object Icon" on page 427

➤ "Has a Non-Matching Child Icon" on page 428

No Matching Object Icon

The **No Matching Object** icon indicates that an object on one side does not match any corresponding object on the other side. The example below shows that the object called **C** has the **No Matching Object** icon because of the change to the disk size (see the **Properties** tab underneath the tree structure on the right side).
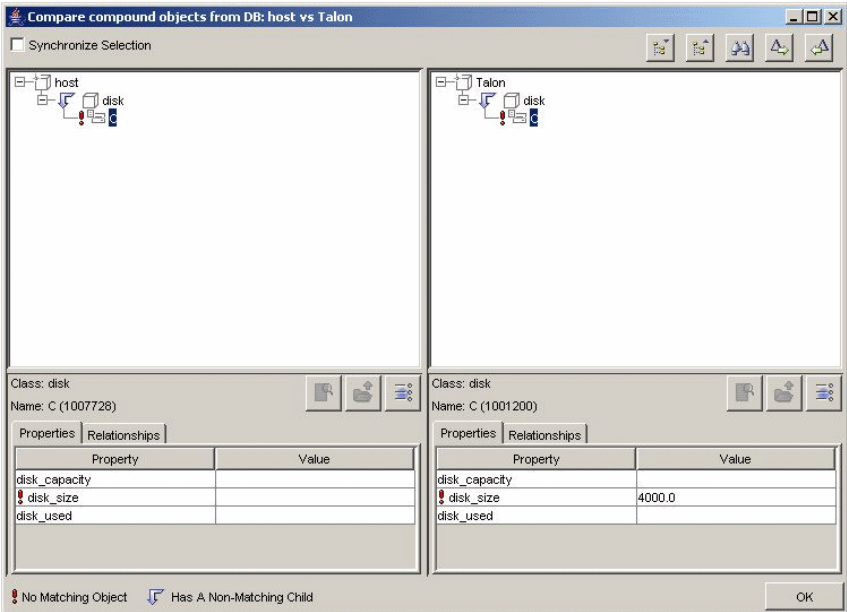
Has a Non-Matching Child Icon

The **Has a Non-Matching Child** icon indicates that the object has a match but one of its children does not have a match. The **disk** object in the illustration above has a **Has a Non-Matching Child** icon due to a change to its child object.

Viewing and Comparing Configuration Files

The Compare compound objects dialog box enables you to view an object's configuration file, and also compare the configuration files of two different objects.

This section includes:

➤ "Viewing an Object's Configuration File" on page 429

➤ "Showing the Difference Between Two Configuration Files" on page 430

Viewing an Object's Configuration File

This section describes how to view an object's configuration file.

To view an object's configuration file:

1 In the Compare compound objects dialog box, select an object of the class **configfile** from the tree.
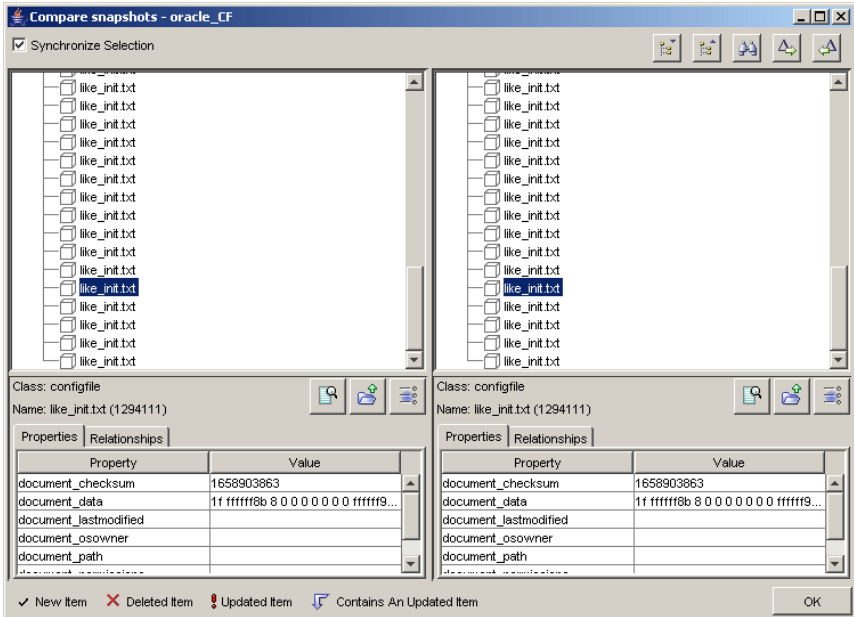
🗁   **2** Click the **Open** button to open the object's configuration file.

Note: This button is only enabled if you have selected an object of the class **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 objects dialog box, select an object of the class **configfile** from the tree on each side.

🗐   **2** Click the **Show Difference** button.

Note: This button is only enabled if you have selected an object of the class **configfile** from the tree.

The differences in the configuration files are marked in blue.

Displaying an Object's Attributes

You can display the selected object's attributes.

To display an object's attributes:

1 In the Compare compound objects dialog box, select an object from one of the trees.

2 Select the **Show Object Attributes** button to open the Object Attributes dialog box.

The Object Attributes dialog box enables you to view the attributes defined for the selected object. The name of the selected object appears at the top. When the **Locked on object** check box is checked, Mercury Application Mapping displays only the attributes of the object in the box at the top. Clear the **Locked on object** check box to display the attributes of any object you select.

Note: This dialog box is for viewing only.

The Object Attributes dialog box has the following tabs:

➤ **Alphabetic** – Displays the object's attributes in alphabetical order.

➤ **Categorized** – Displays the object's attributes according to the class to which it belongs. Click the **Expand** button to display the object attributes for that class.

➤ **Object Attributes History** – Displays a list of attributes in which changes have occurred.

Note: The Object Attributes History tab only displays the attributes whose values have been defined as **Change Monitored** in the Add Attributes dialog box in the Class Browser. For details, see "Creating Classes" on page 374.

The Object 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 Objects Attributes History tab by selecting an attribute from the **Filter by** list.

Note: The Object 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 compound objects dialog box to enable you to select an object on one side and have its corresponding object 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 objects by selecting one on each of the two trees. This allows you to compare two different objects and see their different properties.

The nonmatching properties will be marked by a **No Matching Object** icon in the **Properties** tab (see "No Matching Object Icon" on page 427).

29

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.
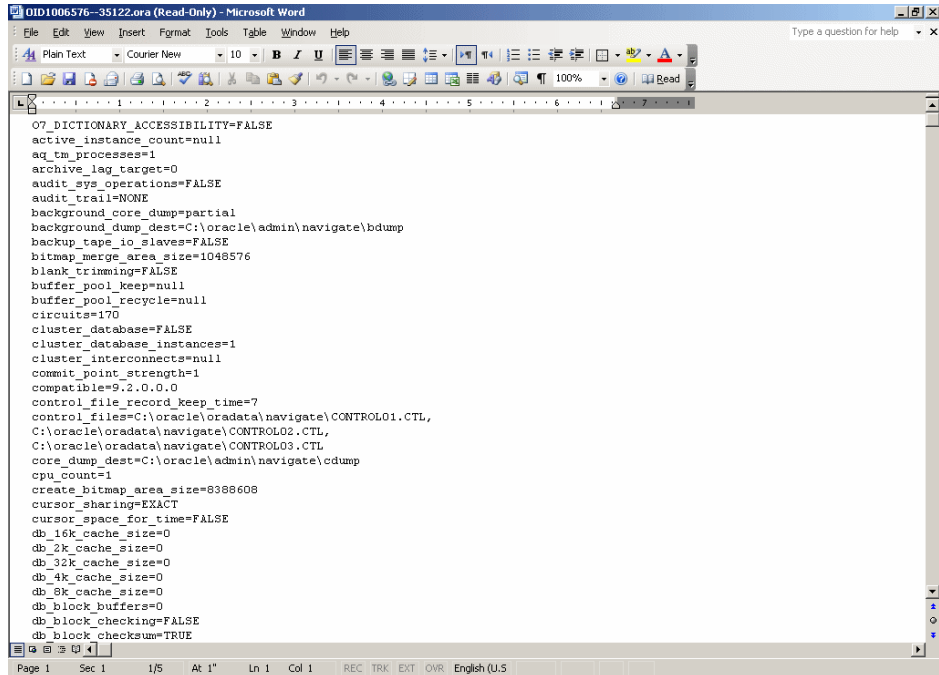
This chapter describes:	On page:
Viewing the Content of a Configuration File	436
Showing the Difference between Two Configuration Files	437
Showing the Differences for the Same Configuration File over Time	438

Viewing the Content of a Configuration File

This section describes how to open a configuration file.

To open a configuration file:

1 Select an object of the class **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.
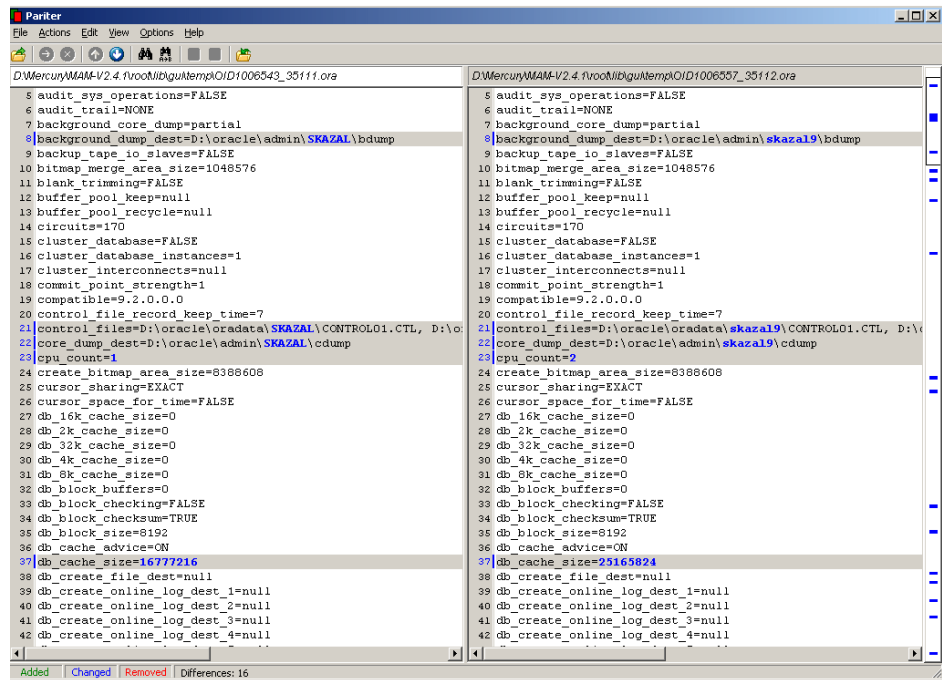
Showing the Differences for the Same Configuration File over Time

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.
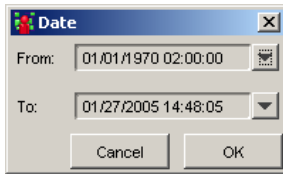
3 Select the dates during which you want to want to display changes that occurred in the file. To do this, do the following:

➤ Click the buttons to the right of the **From** and **To** fields to define the beginning and end of that time period respectively.
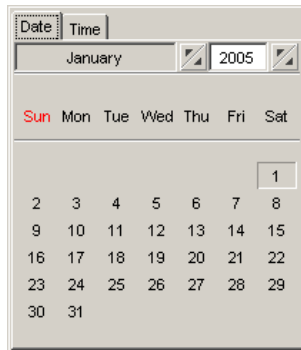
➤ Click the **Date** tab to select the month, date and year.

➤ 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.
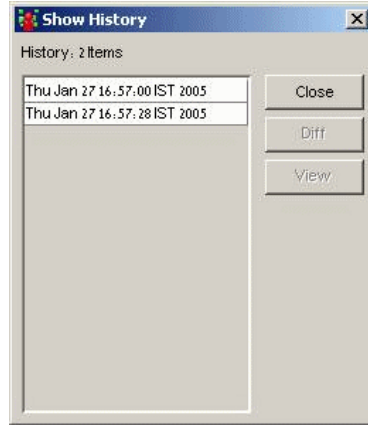
This dialog box lists all the times at which changes occurred in the configuration file.

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.

30

Impact Analysis

This chapter explains how Mercury Application Mapping enables you to assess the impact of infrastructure events.

This chapter describes:	On page:
Evaluating the Impact of Infrastructure Events	441
Modifying a Ticket's Details	443

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 object(s) to the Ticket. The objects you add to the Ticket represent the changes you want to make to the system. Mercury Application Mapping evaluates the impact the objects associated with the Ticket will have on the system when these objects are in downtime. You then generate a report displaying the objects in the system, which will be impacted by the changes specified in the Ticket.

Note: The Ticket Manager can appear in different modes, depending on how you configure the following definitions in the **web.properties** file, located in:

bea81\user_projects\domains\appdomain\appserver\stage_appsdir_mam_war\mam.war\WEB-INF

appilog.web.ticketManagerWrite – determines whether or not you can create new Tickets.

False = does not enable you to create new Tickets. False is the default.

True = enables you to create new Tickets.

appilog.web.ciReadonly – determines whether or not you can modify Ticket properties

False = does not enable you to modify Ticket properties.

True = enables you to modify Ticket properties. True is the default.

To evaluate the impact of infrastructure events:

1 Click **Administration** > **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** button to create a new Ticket. The **Asset Selection** button in the Ticket Details pane becomes a **Create** button.

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** button.

5 Click **Asset Selection** in the Ticket Details pane to open the Ticket Asset Selection screen.

6 In the Object List pane, select the object(s) you want associated with the selected Ticket. Hold the CTRL key down to make multiple selections.

Note: You can also use the Search pane to search for objects and display them in the Object List pane. For details, see "Searching for Objects" on page 202.

---

7 Click the **Add to Ticket** button to move the selected object(s) to the Ticket Details pane.

8 To remove an object from a Ticket, select the object and click the **Remove from Ticket** button.

9 Click **Submit** to associate the object(s) with the selected Ticket.

10 Click the **Generate Report** button to generate a report displaying a list of objects that will be impacted in the system as a result of the simulated changes.

The Change Impact report displays the results as follows:

➤ The objects in the **Impacted Services** tab are grouped according to the hierarchy of the folders and views in the Explorer Tree pane.

➤ The objects in the **Impacted Resources** tab are sorted by class.

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

31

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.

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** – for details, see "Opening a Map of the Currently Selected View" on page 471

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 a user depends on the user's permissions. For details on granting user permissions, see the *Mercury Application Mapping Administration Guide*.

For details on Web browser requirements, as well as minimum requirements to successfully view Mercury Application Mapping, see the *Mercury Application Mapping Installation Guide*.

This section includes:

➤ "Logging In" on page 449

➤ "Logging Out" on page 449

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>/mam/login.html, where **ipaddress** is the IP
address of the Mercury Application Mapping Server and **portnumber** is the
port number of the WebLogic server. The default value is **7001**.

2 Enter the login parameters (login name and password) of a user defined in
the Mercury Application Mapping system, and click **Log In**. 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 superuser change this password immediately to prevent
unauthorized entry. For details on changing the password, see "Changing
Your Password" on page 43.

For details on creating users in the Mercury Application Mapping system,
see 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 of the page.

Mercury Application Mapping WebView Main Window

The Mercury Application Mapping WebView main window includes the following panes:

Search pane   *Explorer Tree pane*   *Refresh button*

Status Bar   *Information pane*   *Map Details pane*
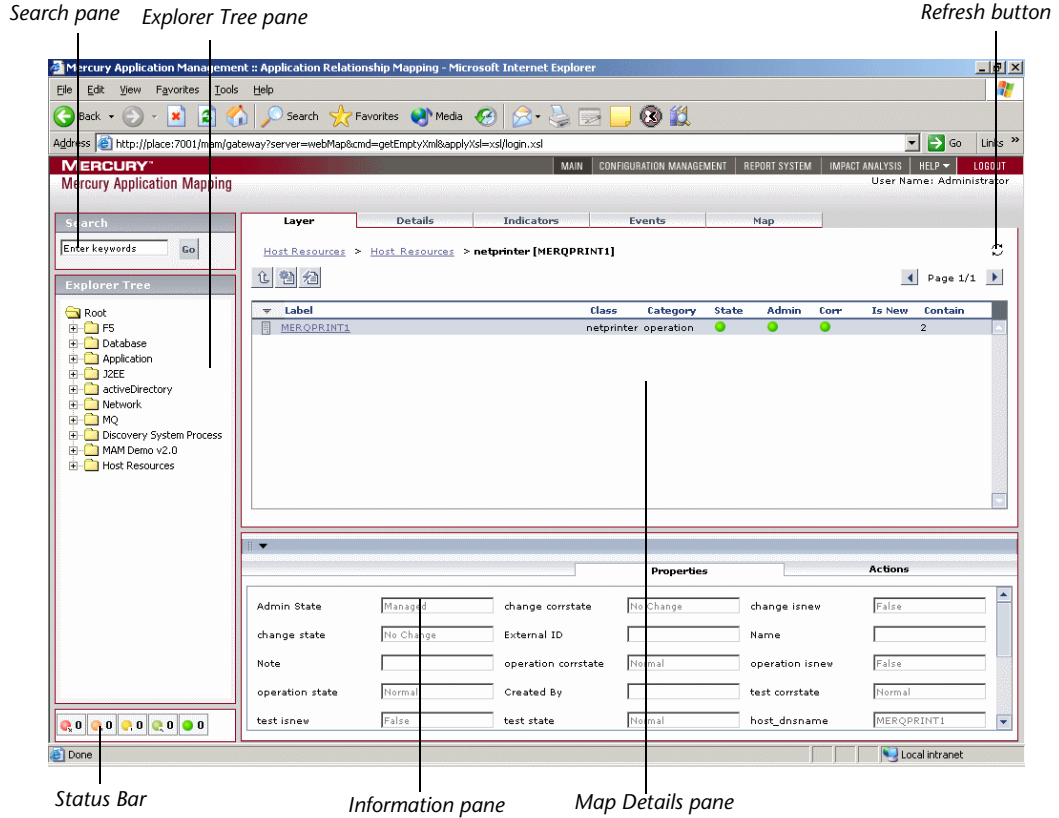
This section includes:

➤ "Search Pane" on page 451
➤ "Explorer Tree Pane" on page 451
➤ "Map Detail Pane" on page 452
➤ "Information Pane" on page 459
➤ "Status Bar" on page 459

Search Pane

You can search for objects that are displayed in the Explorer Tree using the Search pane (for details, see "Searching for an Object" on page 464).

Explorer Tree Pane

The Explorer Tree pane displays a hierarchical tree structure of the map views and the folders in which they are located.

The following table describes the icons in the Explorer Tree pane:

Icon	What it represents
📁	A folder containing a view(s)
🔧	A view

The tree includes folder and view icons and displays three 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 lowest level in the tree hierarchy

Select a level in the Explorer Tree pane to display all objects 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 objects. View is the lowest level in the Explorer Tree hierarchy and does not allow you to drill down to objects or nodes. You can view the view's nodes and objects 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 459).

The Explorer Tree is automatically updated to reflect any change that occurs in the Explorer pane 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. You can drill down to the object level (by comparison, you can view only the root, folder, and view levels in the Explorer Tree). For details on the Explorer Tree pane, see above.

The Map Detail pane includes the following tabs:

➤ Layer tab (for details, see below)

➤ Details tab (for details, see "Details Tab" on page 455)

➤ Indicators tab (for details, see "Indicators Tab" on page 456)

➤ Events tab (for details, see "Events Tab" on page 457)

➤ Map tab (for details, see "Map Tab" on page 458)

Layer Tab

The layer tab displays the topology and status data for the currently selected object.

The objects represented in the map view include topology data as well as basic attributes. This topology data is based on, and produced according to, the way these objects are located and defined in Mercury Application Mapping views.

The Layer tab displays the following details for the selected object:

➤ **Symbol** – The button defined for the class of the selected object.

➤ **Label** – The label of the selected object.

➤ **Class** – The class of the selected object as defined in the class model.

➤ **Category** – The category that is defined for the selected object in the Category Manager dialog box.

Note: If more than one category is defined for the object, the category that has the highest priority as defined in the **Category Priority** list is displayed. For details on setting category priority, see "Defining Map Categories" on page 95.

➤ **State** – The operational state of the object, 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 Status Factor of an Object" on page 99.)

➤ **Admin** – The administrative state of the selected object. Any value other than Managed indicates that any event information coming from the selected object is not reliable because the object is disabled, undergoing testing, or is otherwise non-managed. The Admin State of an object does not propagate upwards to objects at higher levels.

➤ **Corr** – A state that is defined in a correlation rule, and assigned to an affected object once its root cause object fulfills the correlation condition and a correlation event occurs. For details on correlation rules, see Chapter 17, "Defining Correlation Rules."

➤ **Is New** – An exclamation mark in this field indicates that the object'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.

Details Tab

The Details Tab displays the object details according to the format defined in its class report. It is enabled at object level.

The Details tab displays the details of the class report definition of the selected object. The details that appear vary, as they are specific to the class for which you defined the report.

Class reports are defined in the same way you define a host report. For details on defining a class report, see Chapter 22, "Creating System Report Templates."

When creating a class report:

➤ Use the class name as the name of the report. Since report names are case-sensitive, the name of the report must be identical to the name of the class.

➤ Ensure that the class object is first in the node order of the report by verifying that it is at the top of the list in the Set Node Order dialog box.

If a report is not defined for the selected object, Mercury Application Mapping WebView searches for, and displays, a class report for one of the parent objects in the inheritance tree.

Note: Class reports also appear in a separate window when you select the **Class Report** button (for details, see "Mercury Application Mapping WebView Buttons" on page 460).

Indicators Tab

The Indicators tab displays the number of children each object has in the selected view. It is enabled at all levels. It displays information about the objects that exist at the level below the selected object.

The Indicators tab contains the following fields:

➤ **Symbol** – The button defined for the class of the selected object.

➤ **Label** – The label of the selected object.

➤ **Class** – The class of the selected object as defined in the class model.

Note: If an Enumeration is defined for the category, the colors that are assigned for each value in the severity list are also displayed.

➤ **Category** – All categories that are defined for the objects that exist at the level below the selected object.

➤ **adminstate** – A category defined for the view.

➤ **operation** – A category defined for the view.

➤ **No. of Children** – The number of children that exist at the level below the selected object.

Events Tab

You use the Events tab to view event information. The content of the Events tab depends on which button is pressed (**Compound**, **History**, or **QuickFilter is On/Off**). For details, see "Mercury Application Mapping WebView Buttons" on page 460.
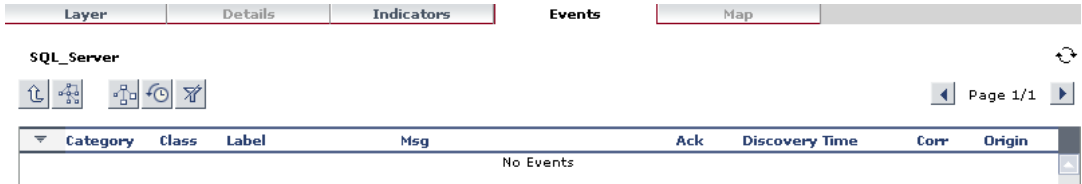
The Events tab contains the following fields:

➤ **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 class name of the event's object.

➤ **Label** – The inclusive label of the object related to the event. Includes other objects that identify the selected object. For example, a disk's label will include 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 object once its root cause object fulfills the correlation condition and a correlation event occurs. For details, see Chapter 10, "Using Map View."

➤ **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 of the currently selected view and navigate between different map levels (from the top level down to the object level) to focus on specific areas of a map. You can also perform Mercury Application Mapping's supported functions. For details on how to perform these functions, see the Mercury, see "Opening a Map of the Currently Selected View" on page 471.
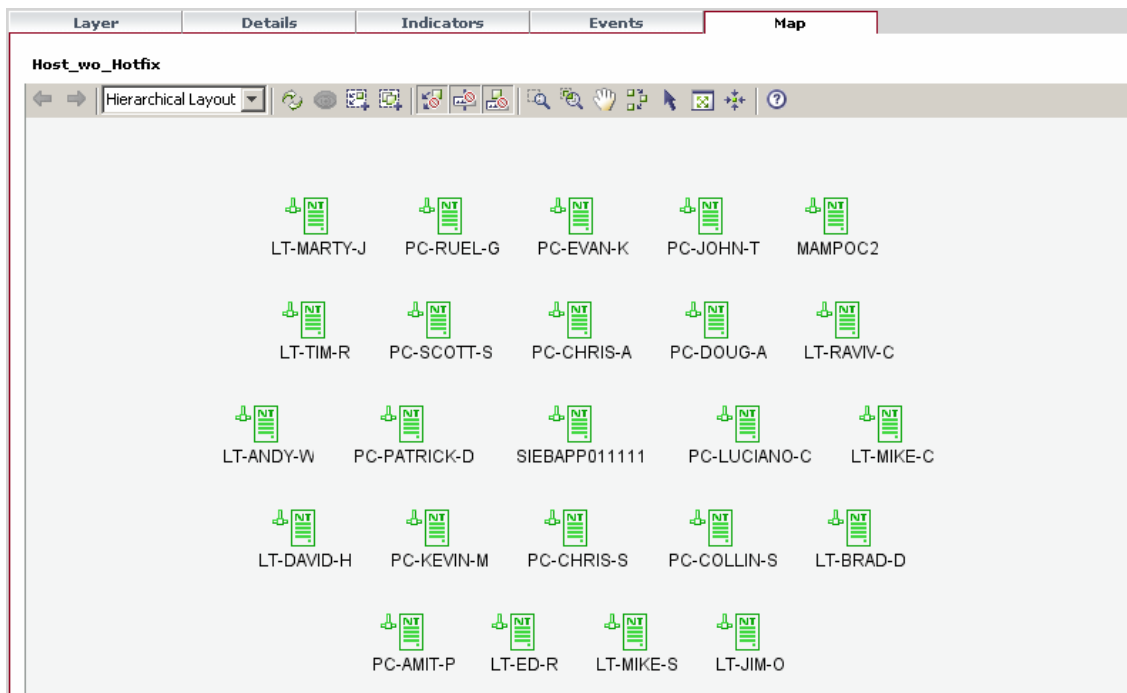
Information Pane

You can view a selected object's properties and you can perform actions, such as showing the root-cause, attaching a note, and so forth.

Status Bar

The Status bar at the bottom of the window displays the number of currently active events for each severity level (critical, major, minor, warning, and normal) (for **Oper State** only).

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 an object.

➤ **Breadcrumbs** – You can return to previous pages within a multi-level context by clicking the appropriate page level. For example, in the following breadcrumbs trail, you would click **Client Server Connections** to return to that page:

➤ **Drill down** – Objects that include child objects are displayed as links. You drill down through the objects by clicking the link.

Mercury Application Mapping WebView Buttons

Mercury Application Mapping WebView includes the following buttons:

Button	Description	Displayed in the Following Tab(s)
	Up in Path – Takes you up in the tree hierarchy. The level that is displayed depends on which map layer is currently displayed: ➤ If the selected view is currently at the View or Folder level, clicking the button takes you up to the Root level. ➤ If the selected view is currently at the object level, clicking the button takes you up through the path one level at a time.	➤ Layer ➤ Events ➤ Indicators
	Show Layer Map – Displays the map of the selected view. For details, see "Opening a Map of the Currently Selected View" on page 471.	➤ Layer ➤ Events
	Class Report – Displays object information according to the object's class report definition. For details, see "Details Tab" on page 455.	Layer
	Event Report – Displays the events for the selected view. For details see "Generating Event Reports" on page 481.	Layer
	Hide/Show Compound Events – Toggles between showing or hiding the active events of the children of the currently selected object. For details, see "Viewing Event History" on page 473.	Events
	Hide/Show Log Events – Toggles between showing or hiding the log events of the children of the currently selected object.	Events

Button	Description	Displayed in the Following Tab(s)
	Quick Filter 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 475. **Note:** This button only toggles if an event filter is defined for the object.	Events
	Quick Filter is Off – Includes the events that are filtered in the Quick Filter dialog box.	Events
	Next Page – Go to next page.	
	Previous Page – Go to previous page.	
	Periodically updates the information with the latest topology and events data.	
	Does not update information.	
	Minimizes and restores the Information pane.	

461

32

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 450.

Searching for an Object

The Search pane enables you to search for objects that are displayed in the Explorer Tree.

To search for an object:

1 In the **Search** box, enter the name of the object or part of it.

2 Click **Go**. An additional tab called **Search Result** appears in the Map Detail 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 Objects Related to a Selected Layer

Select a level in the Explorer Tree pane to display all objects 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 objects. View is the lowest level in the Explorer Tree hierarchy and does not allow you to drill down to objects or nodes. You can view the view's nodes and objects 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 459).

The Explorer Tree is automatically updated to reflect any change that occurs in the Explorer pane in Mercury Application Mapping.

Viewing Details of an Object in the Layer or Events Tab

You can view a selected object's properties and you can perform actions, such as showing the root-cause, attaching a note, and so forth.

To view details of an object in the Layer or Events tab:

Select the object. The Information pane displays the object properties and the actions you can perform on the object:

➤ **Properties** – Displays the attributes of the selected object. The fields that appear vary according to the object you select.

➤ **Actions** – Displays the available actions depending on which tab in the Map Detail pane is selected. For details, see "Using the Actions Tab" on page 466.

Using the Actions Tab

You can use the Actions tab displayed in the Information pane to perform operations on the selected object.

To do that, in the Map Details pane, click the **Layer** tab, select an object (click next to it in the table row) that is affected by a correlation rule and click the **Actions** tab in the Information pane.
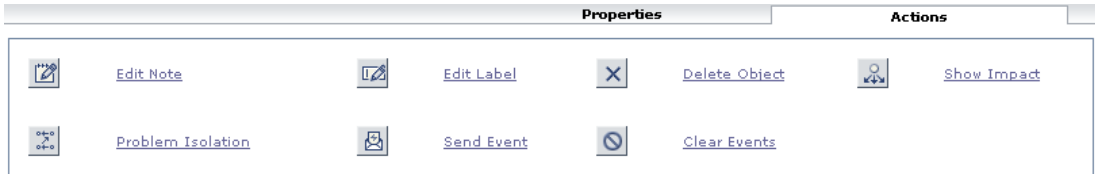
Select one of the actions as follows:

➤ **Edit Note** (for details, see "Editing a Note and Attaching It to an Event" on page 467)

➤ **Edit Label** (for details, see "Editing the Label of a Selected Object" on page 467)

➤ **Delete Object** (for details, see "Deleting an Object" on page 468)

➤ **Show Impact** (for details, see "Showing the Correlation Impact of an Object" on page 468)

➤ **Problem Isolation** (for details, see "Isolating a Problem for an Object" on page 469)

➤ **Send Event** (for details, see "Sending an Event" on page 470)

➤ **Clear Events** (for details, see "Clearing Events from an Object" on page 470

Editing a Note and Attaching It to an Event

You can edit a note and attach it to a node or event. You can use these notes as reminders to yourself or to provide feedback for others.

To edit a note:

1 In the Map Details pane, click the **Layer** tab.

2 In the **Layer** tab, select the object to which you want to add a note.

3 Display the **Actions** tab in the Information pane.

4 Click **Edit Note** to open the Edit Note dialog box.

5 Type the required text in the **Note** box.

6 Click **Save** to save the note.

Editing the Label of a Selected Object

You can edit the label of a selected object.

To edit an object's label:

1 In the Map Details pane, click the **Layer** tab.

2 In the **Layer** tab, select the object whose label you want to edit.

3 Display the **Actions** tab in the Information pane.

4 Click **Edit Label** to open the Edit Label dialog box.

5 Edit the label as required in the **Label** box.

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

Deleting an Object

You can delete an object.

To delete an object:

1 In the Map Details pane, click the **Layer** tab.

2 In the **Layer** tab, select the object to delete.

3 Display the **Actions** tab in the Information pane.

4 Click **Delete Object**. A message is displayed.

5 Click **Delete** to delete the object or **Cancel** to return to the object without deleting it.

Showing the Correlation Impact of an Object

You can select an object which is defined by a correlation rule as a root-cause object, and display all objects that are affected by it in a separate window. For additional information, see Chapter 10, "Using Map View."

To show a correlation impact on an object:

1 In the Map Details pane, click the **Layer** tab.

2 In the **Layer** tab, select a root-cause object.

3 Display the **Actions** tab in the Information pane.

4 Click **Show Impact**.

If several correlation rules are defined for this object, the Root Cause Correlation Rules dialog box appears containing a list of all the correlation rules that define this object as an affected object.

5 Select a correlation rule and click **OK** to open the Impact window.

Note: If only one correlation rule is defined for the selected object, only the Impact window is displayed.

Isolating a Problem for an Object

You can select an object that is affected by a correlation rule and display the root-cause objects that affect it.

To isolate a problem for an object:

1 In the Map Details pane, click the **Layer** tab.

2 Select an object (click next to it in the table row) that is affected by a correlation rule.

3 Display the **Actions** tab in the Information pane.

4 Click **Problem Isolation.**

   If several correlation rules are defined for this object, the dialog box displays a list of the correlation rules that define this object as an affected object.

   If only one correlation rule is defined for the selected object, the Correlation window is displayed immediately.

5 Select a correlation rule and click **OK** to open the Correlation window.
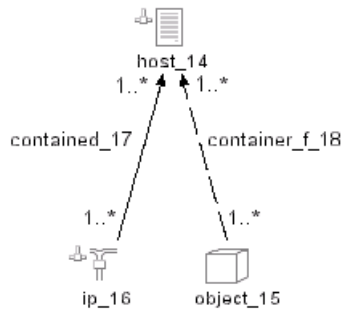
Sending an Event

You can define the category that will be affected by the event to be sent.

To send an event:

1 In the Map Details pane, click the **Layer** tab.

2 In the **Layer** tab, select the object to which you want to send an event.

3 Display the **Actions** tab in the Information pane.

4 Click **Send Event** to open the Send Event dialog box.

5 From the **Category** list, select the category that is affected by the event. The categories (states) **operation, change**, and **test** are system defaults. Other categories (states) in the list are dynamically defined in the State Manager.

6 From the **Enumeration** list, select the required severity.

7 Click **Send Event**. The button that is defined in the Enumeration Manager for the severity you selected is displayed.

Clearing Events from an Object

You can delete an object's events, either those that belong to a certain category (state) or to all categories (states).

To clear events from an object:

1 In the **Layer** tab, select the object from which you want to clear events.

2 Display the **Actions** tab in the Information pane.

3 Click **Clear Events** to open the Clear Events dialog box.

4 Select a category (state) in the **Category** list to delete events belonging to a specific category (state). Note that the list contains categories (states) that are either the system's default categories (states) or are defined in the State Manager.

5 Select **All Categories** to delete events belonging to all categories (states).

Note: When clearing all events for a specific state, the object may still be affected by other state events. Only if you choose to delete the events for all state types, does the object's state return to normal.

6 Click **OK** 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 object level) to focus on specific areas of a map. You can also perform Mercury Application Mapping's supported functions.

To open a map of the currently selected view:

1 In the Map Detail pane, select the required view.

2 Click the **Map** tab. The map of the currently selected view is displayed.
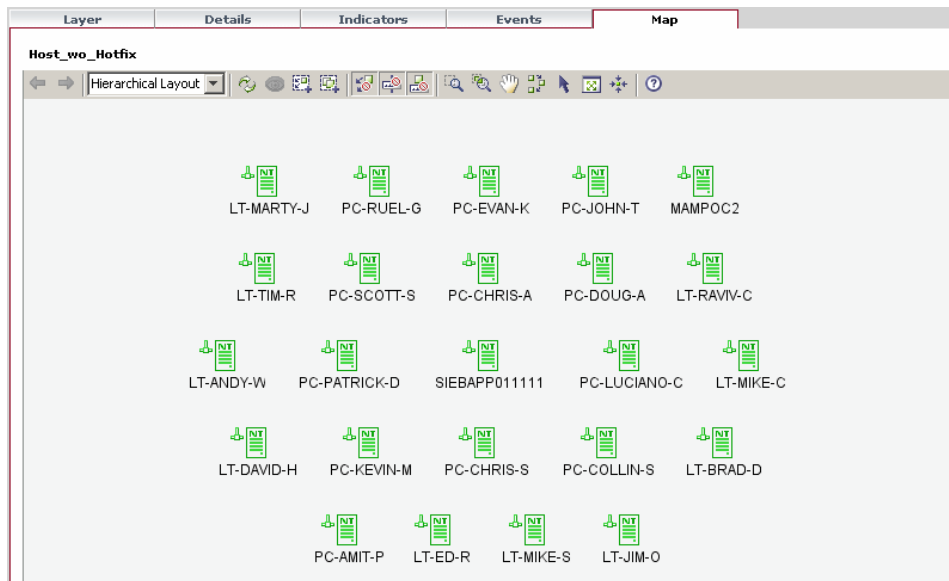
3 Right-click an object to open a menu that enables you perform Mercury Application Mapping's supported functions. For details on the menu options, see "Object Shortcut Menu" on page 140. For details about the buttons that appear in the Map tab, see "Toolbar" on page 34.

Environment Comparison in Mercury Application Mapping WebView

Mercury Application Mapping WebView enables you to compare the snapshots taken of one view at different times. You first take the snapshots in the Mercury Application Mapping main window. Then you can compare the two views in Mercury Application Mapping WebView.

To compare different snapshots of the same view in Mercury Application Mapping WebView:

1 Take different snapshots of the required view in the Mercury Application Mapping main window as described in Chapter 27, "Environment Comparison".

2 In Mercury Application Mapping WebView, click **Configuration Management** at the top.

3 Click the **Compare Snapshots** tab.

4 Select the view in the Explorer Tree pane whose snapshots you took. The Snapshot List pane displays a list of all the different snapshots taken of the selected view. For a description of the fields in the Snapshot List pane, see "Comparing Snapshots Taken at Different Times" on page 396.

5 Select the two snapshots you want to compare and click the **Compare** button to open the Compare snapshots dialog box. For details about the Compare snapshots dialog box, see "Understanding the Compare Snapshots Dialog Box" on page 399.

Evaluating the Impact of Infrastructure Events in Mercury Application Mapping WebView

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 **Impact Analysis** at the top to open the Ticket Manager. For details, see Chapter 30, "Impact Analysis".

Viewing Event Information

You can view information about events as follows:

➤ Display the event history of a selected object (for details, see "Viewing Event History" on page 473).

➤ Find views that contain an object to which a specific event is sent (for details, see "Finding Views that Contain an Event-Related Object" on page 475').

➤ Acknowledge all events relating to a selected object. (for details, see "Acknowledging All Events Related to a Selected Object" on page 475).

➤ Display only specific events in the Events tab for a specific object (for details, see "Filtering the Events You Want to Display" on page 475).

Viewing Event History

You can view the event history of the selected object. Mercury Application Mapping WebView lists the raw and active events that occur in relation to the selected object 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	Hide Compound Events		√	Active events of the currently selected object.
	Hide Log Events		√	
State 2	Show Compound Events	√		Active events of the currently selected object and all its children.
	Hide Log Events		√	
State 3	Hide Compound Events		√	Raw events of the currently selected object.
	Show Log Events	√		
State 4	Show Compound Events	√		Raw events of the currently selected object and all its children.
	Show Log Events	√		

Finding Views that Contain an Event-Related Object

You can find views that contain an object to which a specific event is sent.

To find views containing an event-related object:

1 In the Map Details pane, click the **Events** tab.

2 In the **Events** tab, select an object to which an event has been sent.

3 Display the **Actions** tab in the Information pane.

4 Click **Event Navigator** to open the Event Navigator dialog box.

The dialog box lists the views that contain the object.

5 Select a view and click **Go**. The view appears in the Map Detail pane and the object is highlighted.

Acknowledging All Events Related to a Selected Object

The **Acknowledge** option enables you to acknowledge all events relating to a selected object.

To acknowledge all events related to a selected object:

1 In the Map Details pane, click the **Events** tab.

2 In the **Events** tab, select an object to which events have been sent.

3 Display the **Actions** tab in the Information pane.

4 Click **Acknowledge All**. A red exclamation mark appears in the **Ack** box.

Filtering the Events You Want to Display

You use the Quick Filter dialog box to include only events you want to display in the Events tab for the selected object.

To filter the events you want to display:

1 In the Map Details pane, click the **Events** tab.

2 In the **Events** tab, select an object to which an event has been sent.

3 Display the **Actions** tab in the Information pane.

4 Click **Quick Filter** to open the Quick Filter dialog box.

5 To display events that were sent during a specified period of time:

➤ Click the calendar button 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.

6 Select the required severity levels for the different categories.

7 From the **isAck** list, select an option:

True – display objects whose events have been acknowledged

False – display objects whose events have not been acknowledged

None – display objects whose events are either in an acknowledged or unacknowledged state

8 From the **isCorr** list, select an option:

True – display objects whose events are the result of a correlation rule

False – display objects whose events are not the result of a correlation rule

None – display objects 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.

---

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.

Resizing the Information Tab

You can expand, contract, minimize, or restore the Information pane.

To expand and contract the Information pane:

Move the pointer over the top border of the Information pane until it becomes a double-pointed arrow. Drag the border up or down to expand or contract the pane.

To minimize and restore the Information pane:

Click the arrow icon in the top left corner.

Note: If information is not available for the object, the information pane cannot be maximized.

33

Generating Reports in Mercury Application Mapping WebView

This chapter describes the reports you can generate in Mercury Application Mapping WebView.

This chapter describes:	On page:
Understanding Reports in Mercury Application Mapping WebView	480
Generating Class Reports	480
Generating Event Reports	481
Generating Asset Reports	482
Generating Host Dependency Reports	484
Generating Gold Master Reports	485
Generating System Reports	488

479

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.

➤ **Class Report** (for details, see "Generating Class Reports" on page 480)

➤ **Event Report** (for details, see "Generating Event Reports" on page 481)

➤ **Asset Report** (for details, see "Generating Asset Reports" on page 482)

➤ **Host Dependency Report** (for details, see "Generating Host Dependency Reports" on page 484)

➤ **System Reports** (for details, see Part VIII, "System Reports.")

➤ **Gold Master Report** (for details, see "Generating Gold Master Reports" on page 485)

Generating Class Reports

You can generate a class report for a selected object. The details that appear in the class report vary, depending on how the class report is defined in Mercury Application Mapping.

To generate a class report:

1 Click the **Main** tab at the top.

2 In the Map Details pane, click the **Layer** tab.

3 In the **Explorer Tree** pane, select the required view.

4 Click the **Class Report** button. The class report is displayed.

Note: You can view the same information report when you click the **Details** tab. For details, see "Details Tab" on page 455.

Generating Event Reports

You can generate a report displaying the events for a selected object. You can also define which events should be included in the Event report.

To generate an event report:

1 Click the **Main** tab at the top.

2 In the Map Details pane, click the **Layer** tab.

3 In the **Layer** tab, select an object to which an event has been sent.

4 Click the **Event Report** button to display the Event Report window.

5 To define which events should be included in the report, see the procedure for filtering. For details, see "Filtering the Events You Want to Display" on page 475.

6 Click **OK** to generate the report. The Event Report of the currently selected object is displayed in a separate window. Only the events you have defined are included in the Event Report.

Generating Asset Reports

You can generate a report that lists all the objects in a selected view and their attribute values.

To generate an asset report:

 1 Click the **Configuration Management** tab at the top.

 2 Click the **Asset Report** tab.
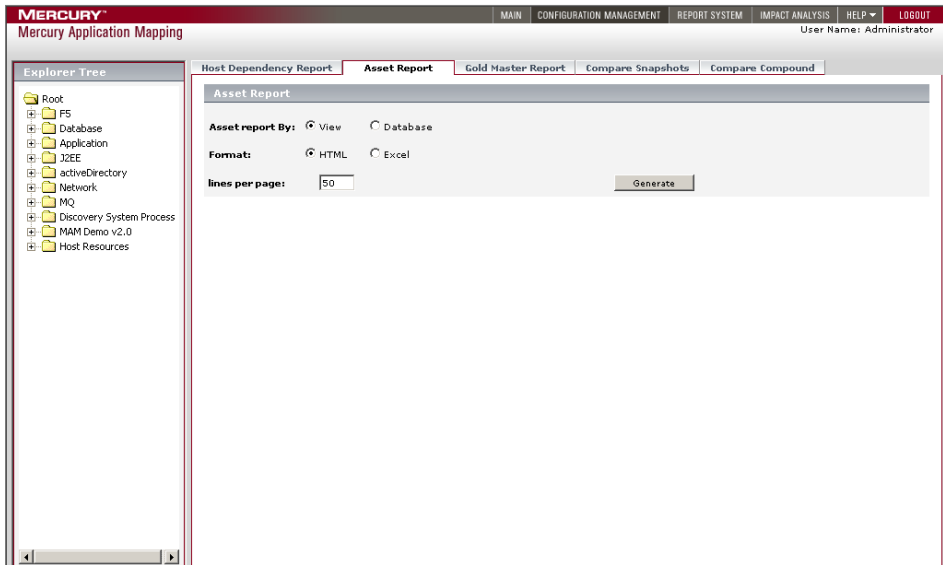
 3 In the Explorer Tree pane, select the view for which you want to create an Asset report.

Note: The report contains only the attributes whose **Asset Data** check box was selected in the **Class Attributes** tab in the Class Browser. For more details, see "Creating Classes" on page 374.

 4 To define what you want the report to include:

➤ Select **By View** to include only objects contained in this view.

➤ Select **By Database** to include additional information from the CMDB. By choosing this option, the report will include objects that are connected by the **container_f** link to the objects in selected view. For example, the report would include the CPU and memory of a server that is contained in the view.

 5 Define the format in which you want to receive the report. You can receive the report in the following formats:

➤ **HTML**

➤ **Excel**

 6 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.

 7 Click the **Generate** button to display the Asset report.

The Asset Report retains the same hierarchical structure as in the Map View.

Generating Host Dependency Reports

You can generate a report displaying the dependencies for a selected object.

To generate a Host Dependency report:

1 Click the **Configuration Management** tab at the top.

2 Click the **Host Dependency Report** tab.
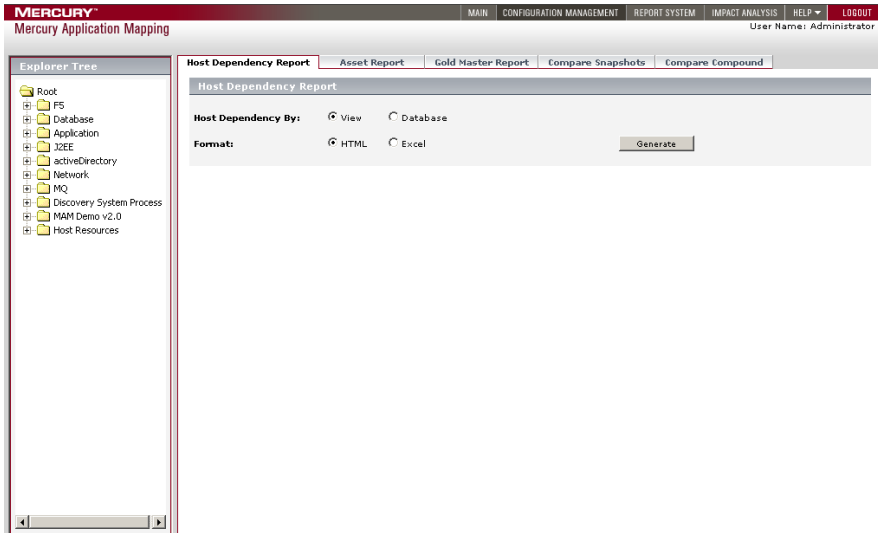
3 In the Explorer Tree, select the view for which you want to create a Dependency Report.

4 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 CMDB that are dependent on the servers in the selected view.

5 Define the format in which you want to receive the report. You can receive the report in the following formats:

➤ **HTML**

➤ **Excel**

6 Click **Generate** to display the dependency report. For details, see "Generating a Dependency Report" on page 266.

Generating Gold Master Reports

Mercury Application Mapping WebView version enables you to generate a report that compares the configuration of a Gold Master object to other objects of the same class. 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 objects to a Gold Master template in the Mercury Application Mapping main window as described in "Comparing the Configuration of Objects to a Gold Master Template" on page 412.

2 In Mercury Application Mapping WebView, click **Configuration Management** at the top.

3 Click the **Gold Master Report** tab to display the following:
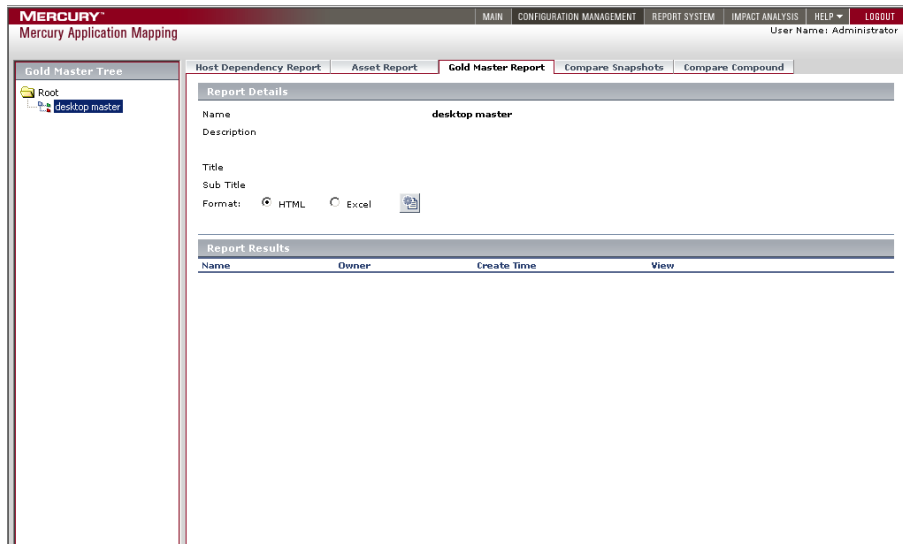
The following table describes the panes contained in the Gold Master Report tab.

Pane	Description
Gold Master Tree	Displays a list of all the Gold Master reports you created 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 To view the results of a report that was generated at specified time intervals, see "Generating a Gold Master Report at Specific Intervals" on page 487.

 4 In the Gold Master Tree pane, select the Gold Master report you want to generate.

 5 In the Report Details pane, select the format in which you want to receive the report:

➤ **HTML**

➤ **Excel**.

 6 Click the **Generate 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 Objects to a Gold Master Template" on page 412.

Generating a Gold Master 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 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 Gold Master report created at a specific time interval:

In the Report Results pane, click the **Generate Report** button in the **View** column of the report you want to generate. The system report is displayed.

Generating System Reports

System reports include information that is defined in Mercury Application Mapping. For details, see Chapter 22, "Creating System Report Templates."

 1 Click the **Report System** tab to display the following:
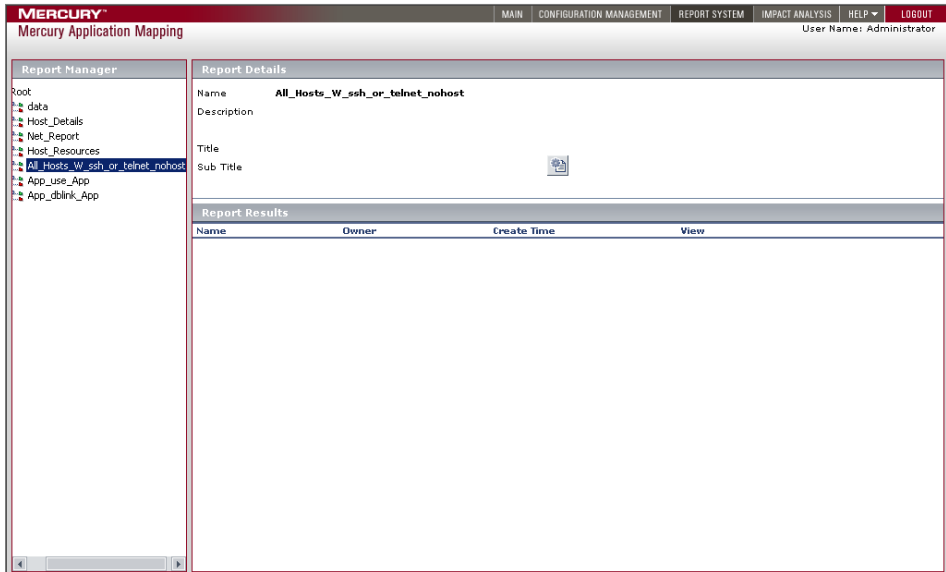
The following table describes the panes contained in the Report System tab.

Pane	Description
Report Manager	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 490.

2 In the Report Manager pane, select the required view.

3 Click the **Generate Report** button to create the report. The report appears in a separate window.

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 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 **Generate Report** button in the **View** column of the report you want to generate. The system report is displayed.

Part XII

Appendixes

A

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 object 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 Map View. 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 object (Disabled, Restricted, Testing, Unknown, Managed, or Unmanaged). Any value other than **Managed** indicates that any event information coming from the selected object is not reliable because the object is disabled, undergoing testing or is otherwise unmanaged. The admin state of an object does not propagate upwards to objects at higher levels.

attribute

> A specific characteristic of a selected object that is defined in the class 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 objects during predefined periods. The availability of an object is measured as a percentage of "up time" versus "total time."

chart

Visual statistics and representation of the network state, which help to identify comparisons, tendencies, and trends in the managed data.

class

An entity contained in the class model.

class model

The set of defined topology objects and their links that describes a managed world.

compound

The children of a selected object.

compound status

Exists in parent objects only and indicates the maximum status of the selected object's children. This status does not include the map status of the parent object.

container link

A functional link that connects a host and an object 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 object once its root-cause object 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 object is always Critical); 2- It can be defined as percentage of the Oper State of its root-cause object (for example, when the correlation event occurs, the Corr State of the affected object is 50% of the root-cause object's Oper State.) When several correlation events are simultaneously related to an object, the correlation state is the maximum severity of all the events.

discovery

> The process of finding resources within an IT infrastructure.

Discovery Manager

> The Discovery Manager is 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 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 object and any objects derived from the selected object (inheritance objects) are displayed in the Map View.

Explorer pane

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 object

An object created from a raw event that contains all the attributes of the event, including the object to which it is related.

event severity

An attribute of an event object 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 object.

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 link

Classes that inherit the attributes and behavior (such as the label function) of the class above it. For example, the IPserver and IPclient classes inherit attributes from the IPport class.

interdependencies

See the glossary item **link**.

join_f link

A link that connects two objects 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 an object 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 objects are divided into layers and you can drill down from one layer to the next to see more objects.

link

Also known as interdependencies. An entity that defines the relationship between two objects. Links can describe either a physical or logical connection between two objects, 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 Map View 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 object, 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 object's map status is **5** (5*100/100). If the weight is defined as 60%, the object's status is **3** (5*60/100)**.** If the object has children, the map status value is the maximum status of the object and all its children.

object

An entity in a managed world that is distinguished from other objects by its class and key attributes.

oper state

The current operational state of the selected object, which is the maximum severity of any event currently occurring to the object.

organizational rules

Rules that define how linked objects are displayed in the Map View. 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.

Probe Gateway

> A communication server located at a remote site. The server connects Mercury Application Mapping and the remote site, and provides a way of communication between the Discovery Manager and one or more Probe Managers. The Probe Gateway handles probe management, communicates on its behalf with the Discovery Manager, and increases its functionality.

Probe Manager

> A component that hosts discovery methods and provides management services that enable their manipulation. The Probe Manager receives a task, delivers its parameters to a discovery method, manages the discovery method performance, and returns the task results to the Probe Gateway.

propagation

> The transmission to higher view levels of managed object data, such as status or blink.

raw event

> A message that notifies Mercury Application Mapping that a specific change has occurred in a managed object.

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 object that is related to the event.

status factor

A value (in percent) that defines the significance of the object status. This is useful for reducing the significance of objects whose status can fluctuate on a regular basis, such as processes, which continuously change. By reducing the significance of such objects 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 object 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 objects and links that can be created without any prerequisites, that is, the grouped objects 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 links between objects and create visual images of IT infrastructure resources.

topology

(From the Greek *topos* meaning "place.") A visual description of the configuration or arrangement of an IT infrastructure, including its nodes and connecting links.

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 objects and links is deposited, grouped, and updated in the form of class definitions according to object-oriented methodology.

topology object

> Any member of a class that exists within the topology of the managed world.

view

> A collection of objects and links represented by icons. These objects and links are the result of a TQL query or an instance view, and are displayed as a view according to display and organizational rules that are assigned to them. Each object/link can be presented in multiple views or by multiple icons in different view layers.

view indicators

> A set of pre-defined, dynamically changing diagrams, which contain statistical data about several significant attributes of the system's views.

viewing layer

> A two-dimensional arrangement of topology objects to be viewed by a user.

Appendixes

Index

T