BrightStor® Resource Manager

Windows Client Guide

6.2

Computer Associates™

Contents

Chapter 1: Introduction

Chapter 2: Overview

Chapter 3: Installation and Configuration

Chapter 4: Navigating the Windows Interface

Index

Chapter

1

Introduction

The *BrightStor Resource Manager Windows Client Guide* provides instructions for the operation of the BrightStor Resource Manager Windows Client. It is a companion to the *BrightStor Resource Manager User Guide* and the *BrightStor CA-Vantage Storage Resource Manager Guide*.

The Windows Client enhances your storage management abilities in the following areas:

- Analyzing storage
- Planning and forecasting
- Defining constructs
- Automating policies
- Measuring effectiveness
- Modeling changes
- Performing ad hoc operations
- Troubleshooting
- Recovering from hardware errors
- Scheduling storage management tasks
- Managing storage on distributed networks

You can perform storage administration from the Windows Client graphic user interface. You can monitor volumes, errors, and other objects related to storage management. Each object of interest can be filtered and sorted using combinations of any of its attributes. For example, the data sets in each storage group, application group, or even the entire system, can be filtered, sorted and viewed as if they resided on a single virtual volume. The online facilities are complemented by logging options and historical analysis reports for any time period you want to examine.

Using generalized threshold monitors for objects, you can automate responses to various conditions, as well as schedule the time of the response. You can implement a system-managed storage environment, in which BrightStor Resource Manager automatically takes the actions you specified, or you can create a manual environment in which BrightStor Resource Manager issues messages to the operators when those thresholds are reached.

In addition to integrating functions of several storage management products, BrightStor Resource Manager also has the ability to automate the functions of other vendor products. An effective implementation of BrightStor Resource Manager and its components corrects storage management concerns before they become problems, increases disk utilization, and reduces both daily operating costs and the need for new investment.

Intended Audience

The information herein is intended for data center storage administrators and system programmers who install and use BrightStor Resource Manager. Readers of the manual are expected to be familiar with the Microsoft Windows operating environment.

Synopsis of the Manual

This document is organized into the following chapters:

Chapter 1, *Introduction,* provides a general explanation of the content and conventions of the Windows Client Guide.

Chapter 2, *Overview,* provides the overview of the BrightStor Resource Manager system components, available objects, and services.

Chapter 3, *Installation and Configuration,* contains step-by-step instructions on how to install and customize the BrightStor Resource Manager Windows client.

Chapter 4, *Navigating the Windows Interface*, explains how to use BrightStor Resource Manager to view objects and manipulate BrightStor Resource Manager object tables.

Conventions

The following typographical conventions are used:

Names of tables and dialog boxes are capitalized. For example, volumes is capitalized in the following sentence because it is a table.

The Volumes table contains the following fields.

Menu branches and commands, boxes and buttons in dialog boxes, and icon names are represented in bold type. For example:

Select **Function Status** from the **View** menu.

Click the **Graph** button under the menu bar.

Click the **Refresh** icon.

Steps and instructions are preceded by a number. For example:

1. Identify the member to be submitted.
2. Specify a time.

Single-step instructions are designated by a checkmark:

✓ Select **Statistics** from the **View** menu.

When you enter data in a dialog box, or when BrightStor Resource Manager displays a message, the text is reproduced in a bold, mono-space font, as shown below:

```
Printed &pr out of &tt lines
```

Revision Information

This manual describes the operation of BrightStor Resource Manager Release 6.2.

Related Publications

The following publications contain useful information that may expand your understanding of BrightStor Resource Manager:

BrightStor Resource Manager Related Publications

For z/OS:

BrightStor CA-Vantage Storage Resource Manager Getting Started

BrightStor CA-Vantage Storage Resource Manager User Guide

BrightStor CA-Vantage Storage Resource Manager Messages Guide

BrightStor CA-Vantage Storage Resource Manager Reference Guide

For Open Systems:

BrightStor Resource Manager User Guide

BrightStor Resource Manager Getting Started

Related System Publications

CA-Disk Installation Guide

CA-Disk Systems Guide

CA-Disk User's Guide

CA-Allocate User's Guide

Other Publications for z/OS

ES2 User Modifications and Macros

IBM: VTAM Programming

IBM: SMP/E Reference Manual

IBM: MVS/ESA JES2 Customization

IBM: MVS/ESA Installation Exits

IBM: VTAM - Resource Definition Reference

IBM: MVS/ESA System Management Facilities (SMF)

IBM: DFSMShsm Storage Administration Reference

IBM: DFSMShsm Managing Your Own Data

IBM: DFSMSrmm Guide and Reference Application Programmers Interface

Microsoft: Windows User's Guide

StorageTek: HSC System Programmer Guide for MVS

Computer Associates: CA-1 Systems Programmer Guide

Computer Associates: CA-Dynam/TLMS System Programmers Guide

Computer Associates: AutoMedia User Guide

Chapter 2

Overview

BrightStor Resource Manager is a storage management system for Open Systems platforms and MVS environments.

The main function of BrightStor Resource Manager is to provide a solution for storage management problems. Its services are designed to simplify and automate both personal and system-managed tasks through a unique combination of viewing, analysis, and automation facilities.

The Windows Client is a graphic user interface running under Microsoft Windows on a PC workstation. The Windows Client brings to the PC user the complete BrightStor Resource Manager functionality. The Windows client communicates with the z/OS host and Open Systems Application Server through TCP/IP. Users can request system and storage management status information and initiate actions from the Windows client. Storage management data collected, stored, and filtered on the z/OS mainframe or Open Systems Application Server is displayed at the Windows client taking advantage of the graphic capabilities of the user interface. The Windows client provides enhanced editing and printing facilities, and allows data to be exported to other Windows applications for further analysis and reporting.

This chapter provides the overview of the BrightStor Resource Manager system components, available objects, and services.

Objects

The BrightStor Resource Manager object-oriented approach is designed to automate both personal and system-managed storage management tasks through a range of viewing, analysis, and automation facilities.

Storage management operations can be performed manually as follows:

1. Select an object such as a volume, a storage group, a data set group, or a data set.
2. On the selected object, perform a set of actions that can be selected from the applicable menus or initiated by means of commands. These are referred to as *ad hoc* operations.

To automate many storage management tasks, you can define services (on Open Systems) or scripts (on z/OS) and take automatic action when certain conditions within the system are met. There are two main types of automation scripts: object automation and message automation. Message automation is event-driven; object automation is timer-driven. While the details of implementation of the two differ, the general principles are the same. Wizards operating under the Windows client facilitate the definition of both types of scripts.

Additionally, you can define thresholds for various space and fragmentation conditions and specify the job to be performed when these thresholds are exceeded. The system automatically submits the jobs when the thresholds are exceeded.

BrightStor Resource Manager provides information about its own activity through several objects, as described in the following sections.

Message Log Object

The message log object contains messages about the normal stages of an operation, suspect or error conditions, and diagnostic information regarding a detectable BrightStor Resource Manager or related software malfunction.

The activity log server writes all messages issued by BrightStor Resource Manager to a central file called the Activity Log. When BrightStor Resource Manager is started, it attaches to the activity log server. The server's initialization routine creates an in-memory queue. Other tasks in the system place messages in the queue for processing by the activity log server.

The activity log server allocates and opens the Activity Log when the first message appears in the queue. If no generations of the Activity Log exist, the activity log server creates one. When the current generation is full, the subtask creates a new generation.

Work Scheduler Object

The scheduler object lists all scheduled events, both automatic and user initiated. Jobs are scheduled both manually (by the storage administrator) and automatically (by BrightStor Resource Manager) in response to the violation of some condition.

The Work Scheduler provides a consistent set of scheduling services for all job submissions within the BrightStor Resource Manager environment, for both *ad hoc* and automated submissions for volumes, storage groups, data set groups, and data sets. Events can be scheduled by an absolute time specified as a day, hour, and minute, or as a delay specified in years, days, hours, and minutes. The services are called from other components whenever they are needed.

Logged-On Users Object

The logged-on users object shows a list of users currently logged into BrightStor Resource Manager.

Mail Object

The mail object lists mail messages sent to the logged in BrightStor Resource Manager user's mailbox, including the text of each message, the sender, the receiver, and the date and time the message was issued.

The BrightStor Resource Manager base system provides services for sending mail to BrightStor Resource Manager users. SEND commands can be issued from the operator console, from the Windows and View 3270 interfaces, or by using a batch utility program.

When the base system is started, a Mailbox Server is also attached. The server is used to store messages for users who are not logged on, allowing the messages to be retrieved at a later time. The server is also responsible for saving the messages in a checkpoint data set when BrightStor Resource Manager (or just the server) is shut down, and reloading them when BrightStor Resource Manager (or the server) is restarted. The server also deletes old messages according to a user-specified parameter.

Messages can be sent to a specific user ID or broadcast to all users. Use the mail service as part of BrightStor Resource Manager shutdown procedures to notify users of the shutdown. Similarly, the batch mail utility can be added to automation jobs to notify the staff when needed.

Automation Manager (z/OS only)

The BrightStor Resource Manager Automation Manager provides the following for z/OS objects and systems:

Script-based automation - use this feature to create scripts that allow BrightStor Resource Manager to take corrective action automatically in response to various types of system messages and events in the system.

Scripts fall into two main categories

- Message automation used to initiate corrective action by monitoring system console, DFSMShsm, or CA-Disk messages. The scripts collect, filter, and accumulate or count the messages. When conditions defined in the script are met, the script triggers the submission of the appropriate job.

- Object automation used to monitor any BrightStor Resource Manager object and initiate corrective action.

  Scripts are a convenient method of taking automatic action when certain conditions within the system are met. Two main types of automation scripts support two types of automation schemes: object automation and message automation. Message automation is event-driven; object automation is timer-driven. While the details of implementation of the two differ, the general principles are the same. Wizards operating under the Windows client facilitate the definition of both types of scripts. The resulting objects can be viewed both from the Windows and from the View 3270 interfaces.

  Automation scripts are stored in PDS members from which they are loaded into memory and activated. Several scripts can be loaded and active at the same time.

Threshold automation - use this feature to define thresholds for volumes, storage groups, and data set groups and instruct BrightStor Resource Manager how to respond automatically to violations. This feature supports:

- Data collection for volume space and fragmentation as well as for storage group space.
- Construct definition and automation for VTOC space, for volume space and fragmentation, for storage group space, and for data set group space.

Script Language

BrightStor Resource Manager provides a script language to help you write command statements that correspond to "point-and-click" actions. For example, you can use a script to select an object, set filter and sort criteria, execute the request to get the records, write the records to log files, substitute variables from the records into model JCL templates, and schedule or submit the resulting job.

Additional features allow you to specify the time a script should be enabled and disabled; the maximum number of times the script can be executed during its enabled period; and whether selected records should be counted or accumulated for processing.

The BrightStor Resource Manager Windows interface provides services to generate and save several types of user-defined scripts such as message automation scripts, general object automation scripts, and object logging scripts. The scripts are saved in partitioned data sets. A system script is another type of script used internally by BrightStor Resource Manager. System scripts are predefined and included as part of the installation process. If automation or logging scripts are defined, system parameters can be set to start the corresponding script servers automatically.

Script servers are responsible for loading the scripts from the PDS and dispatching them appropriately. The Message Automation Server and the Log Server are each described briefly in sections that follow.

Logged and Captured Objects

BrightStor Resource Manager collects trend data into log data sets using logging and capturing. The resulting log data is used primarily by the storage administrator. Captured data intended for individual use is stored under different naming conventions, so as not to overwrite the log data managed by the storage administrator.

The Windows interface provides a log operations wizard for generating scripts to log BrightStor Resource Manager objects to sequential files on disk. In general, you select an object from a list, specify when or how often it is to be logged, and save the definition. System parameters provide high level indexes for the log data sets and control parameters to indicate when to append to an existing log and when to start a new one.

When an object is logged, its data is written to the object's own log file.

Trend Report Objects

The Trend Reporter is an auxiliary to the BrightStor Resource Manager Windows client. Trend report objects let you view the history of selected storage management objects by graphically presenting trend data accumulated on the host in the BrightStor Resource Manager MVS log files. The trend reporter extracts data about selected BrightStor Resource Manager MVS objects over time and stores it on the Windows client. Based on this data, the program generates reports that graphically display the changes that the selected objects underwent during the collection period.

The trend reporter can also perform an analysis of the accumulated trend data and project it into the future for any period of time. This feature is especially important for capacity planning.

The log records can be written to the MVS SMF files, or to sequential data sets on disk, managed by the server.

Bringing z/OS Host Log Data to the Windows Client Trend

The Scheduler component of the Trend Reports program periodically collects selected data from the BrightStor Resource Manager log files on the z/OS host, and stores it with the trend object.

The host log file contains a complete snapshot of the object for which the data is collected: all the records and all the fields. The Scheduler can be instructed to filter this data and copy only selected fields and records.

Chapter

3

Installation and Configuration

This chapter gives important information to consider prior to installation and gives instructions on customizing your new installation.

Installation

Refer to the *BrightStor Resource Manager Getting Started* guide for instructions on installing the BrightStor Resource Manager Windows Client.

Uninstalling BrightStor Resource Manager

Navigate to Add/Remove Programs tab of the Windows Control Panel, select BrightStor Storage Manager, and click Remove. The uninstall program also removes the **TD.MDB** file that contains the user definitions. When this file is removed, all user definitions are removed and, after a new installation, the system starts with default definitions.

Note (z/OS only): If you are updating or upgrading from a previous version to BrightStor Resource Manager 6.2 and use the automatic uninstall procedure that is part of Setup, the old **TD.MDB** file that contains your configuration information is copied to a new location in a BrightStor Resource Manager 6.2 directory. If, however, you uninstall BrightStor Resource Manager manually, a new **TD.MDB** is created for BrightStor Resource Manager 6.2 with default definitions. In this case, to preserve user definitions you must save a copy of **TD.MDB** in a separate location before you uninstall then copy it manually to the appropriate BrightStor Resource Manager 6.2 directory after you have completed the new installation.

Client Customization

BrightStor Resource Manager client customization involves the following steps:

- Starting BrightStor Resource Manager
- Defining one or more hosts with which you want the Windows client to communicate (z/OS only)
- Setting the activation parameters of the BrightStor Resource Manager client
- Optionally, defining object colors
- Optionally, defining user views
- Optionally, defining working sets
- Optionally, defining audit and system scripts (z/OS only)

Starting BrightStor Resource Manager Windows Client

To start BrightStor Resource Manager Windows Client:

✓ Click on the **Windows Client** icon in the BrightStor Resource Manager program group, or click **Start**, then **Programs**, then **BrightStor Resource Manager**.

Note: If you installed the Open Systems option, the Application Server must be running before you start the Windows Client.

If you start BrightStor Resource Manager with the Open Systems option for the first time after a new installation, when the Login dialog appears, enter ADMIN in the **User Name** text box and click **OK** without supplying a password. You are now logged into the BrightStor Resource Manager Windows Client as the default storage manager. Refer to the online help topic Administrator Utility for information on changing your password.

When you start BrightStor Resource Manager for the first time after a new installation or upgrade, the host list and main object tree are activated in addition to the main toolbar and menu. Subsequently, your settings in the **Activation** dialog of the **Tools** menu determine the components activated when you start BrightStor Resource Manager (refer to Setting Activation Parameters).

To start the Windows Client with the default parameters, regardless of the activation parameters defined in the Activation dialog as follows:

1. Right-click on the Windows Client icon and select **Properties**.
2. When the Shortcut to vantage.exe Properties dialog appears, click on the **Shortcut** tab and add the **/N** switch to the expression in the **Target** field.
3. Click Apply and close the dialog box, then start the Windows Client by clicking its icon.

z/OS only: To start the Windows Client with a database other than the default (**TD.MDB**), without loading the hosts, user views, and other objects defined by the user during previous activations of the program:

1. Right click on the Windows Client icon and select **Properties**.
2. When the Shortcut to vantage.exe Properties dialog appears, click on the **Shortcut** tab and add the **/F filename** switch to the expression in the **Target** field. The file name you supply replaces the **TD.MDB** file, the default database that contains the user definitions.
3. Click **Apply** and close the dialog box, then start Windows Client by clicking its icon.

Behavior of the TCP Server

The TCP Server provides a connection to the host both to Windows Client and to the Trend Reporter. The login process is controlled by the TCP Server, so you can log in to both programs with the same User ID and password.

If you exit the Windows Client but leave the Trend Reporter running (as, for example, when the Trend Reporter is engaged in a long data collection process), the TCP Server remains active. If you start the Windows Client again while the Trend Reporter is still active, the program does not ask you to log in. Make sure that this condition does not lead to a security breach. See the Scheduler Security section of this guide for more details.

Defining a Host

Operations performed on BrightStor Resource Manager objects always refer to one or more z/OS hosts, or a single Open Systems host. Each z/OS host has a unique IP address and port number. To log on to the z/OS host, a user ID and password are necessary.

Note: Both the user ID and password are sent to the host in encrypted form.

The Host List window lists the hosts that have been defined on the Windows Client. (Setup does not erase existing host definitions, so if you have defined hosts with a previous version of BrightStor Resource Manager they now appear in the Host List window.)

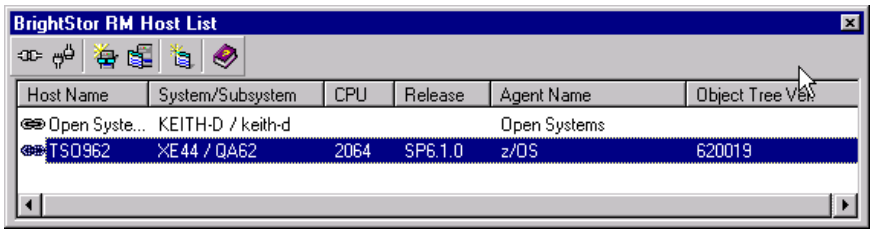

Hosts can have the following statuses:

- Connected
- Not connected
- Connecting

Each status is identified by its characteristic icon.

The Open Systems host (the Application Server), is defined and connected by default after the Client installation. For z/OS hosts, you can use either the host definition wizard or the definition form to add a new host. The following procedure describes the use of the wizard.

1. Right-click in the Host List window and choose **New Host**, then **By Wizard** from the pop-up menu.
2. Enter a name for the new host.
3. Select an agent running on the host from the **Agent Name** drop-down list.

   If you selected an agent with a direct connection mode, click **Next**. If you selected an agent with an indirect connection mode, click **Finish**.
4. If you selected an agent with a direct connection mode, enter the IP address and port number of the system to which you want to connect, then click **Next**.

   If you leave these fields empty, the host fills them when the connection is made.
5. Optionally, enter your user ID and password on the mainframe to which you want to connect. Reenter the password to confirm. Click **Next**.

   If you leave the user ID and password fields empty, each time you connect to the host, BrightStor Resource Manager displays a login dialog asking for your user ID and password.

   If you enter a user ID but no password during host definition, each time you connect to the host, BrightStor Resource Manager displays a login dialog asking for the password.
6. If for debugging purposes you need to create a trace on the host for messages exchanged between the host and the client, enter the value EWS10T in the **Host Trace** field. Do so only if explicitly instructed to by the technical support staff.
7. Optionally, check the **PC Communication Trace** box to save messages relating to the host-client communication in a log file on your workstation. Click **Next**.

8. Click **Test** to verify the IP address and port number you specified.

   a. If the test is successful, BrightStor Resource Manager displays the message:

      The connection has been tested and found to be working.

   b. If BrightStor Resource Manager cannot establish a connection with the host, it displays the message:

      Connectivity test failed.

      Verify the IP address and port number, and try again.

   After the connection tests successfully:

9. Click Finish to complete the definition.

This ends the action of the wizard. The newly created host is added to the Host List.

To review the definition and make changes if necessary, select **Host definition** from the right mouse button menu. The host definition form shows all your parameter settings.

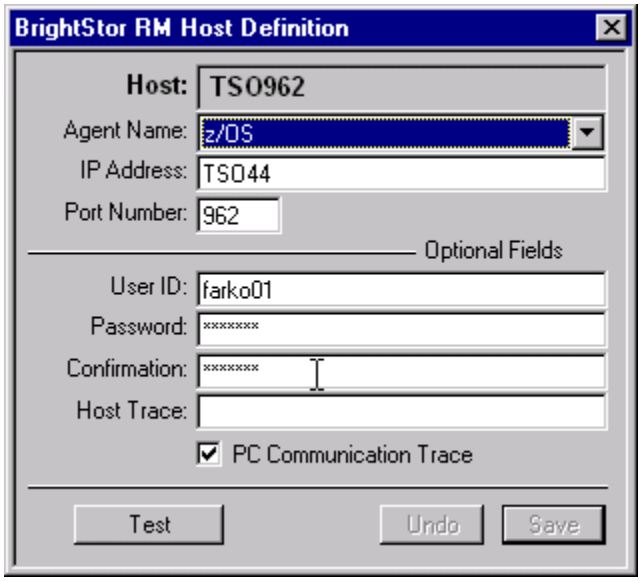

You can modify any of the parameters, except the host name. Click **Save** to apply the host definition changes you make on this form.

Setting Activation Parameters

After you successfully install and activate BrightStor Resource Manager, you can specify the manner of its future activation. To do so:

✓ Select **Activation** from the **Tools** menu of BrightStor Resource Manager.

Use this form to specify the BrightStor Resource Manager features and components you want active the next time the system is started. Information in the form is organized into several tabs: **General**, **Hosts**, **User Views**, **Working Sets**, and **Consoles**.

After you make your selections on all the tabs, click **OK**. Your choices become active the next time the system is started.

General Tab

Use the **General Tab** to specify the windows you want open when BrightStor Resource Manager is started. You can choose to have the following windows open at start-up:

- Object Tree
- Host List
- Active Objects
- Console Tree

If you check the **Run in Demo mode** box, BrightStor Resource Manager starts without connecting to an z/OS host. Open Systems must always be connected to its host. If any tables have been saved with the **Save Data for Demo** function (see section Save Data for Demo), these tables can be accessed in demo mode.

Use the **Data Selection Method in Tables** radio buttons to determine the method of cell and row selection in BrightStor Resource Manager tables. In **Block Mode**, you can select individual cells or click the mouse and drag it to select any rectangular block of cells. In **Line Mode**, you can click the mouse to select individual rows (lines), **Shift**-click to select a group of contiguous rows, or **Ctrl**-click to select non-contiguous rows.

If you check the **Automatic Table Column Adjustment** box, the width of all the columns in all the tables is automatically adjusted to the width of the data.

If you check the **Start Scheduler** box, the BrightStor Resource Manager Scheduler is started together with the client software.

If you check the **Inhibit Scheduler...** box, the Scheduler does not break the client connection to the host after completing a scheduled task, even if no other objects are active on that host.

Hosts Tab

Use the **Hosts** tab to specify the host connections you want to activate when BrightStor Resource Manager is started.

User Views Tab

Use the **User Views** tab to specify the user views you want to activate when BrightStor Resource Manager is started. Available views are listed in the top pane; selected views in the bottom pane. You can move views between the two panes by selecting them and clicking the up and down arrows.

Working Sets Tab

Use the **Working Sets** tab to specify the working sets you want to activate when BrightStor Resource Manager is started. Check the box next to each working set you want to activate.

Consoles Tab

Use the **Consoles** tab to specify the consoles you want to activate when BrightStor Resource Manager is started. Check the box next to each console you want to activate. See also the section Consoles (z/OS only) for more information on consoles.

User Customization

BrightStor Resource Manager stores all user definitions (user views, host definitions, working sets, consoles, and so on) in a database. By default, the database is named **TD.MDB** for z/OS hosts, or TD_XXXXX.MDB (where XXXXX represents the userid used to login to the Windows Client) for Open Systems hosts or for z/OS users who elect to work with the login option.

For each user you define, BrightStor Resource Manager creates a new database file named **TD_xxx.MDB**, where **xxx** represents the user ID. Each user can create and save a custom system definition, with personalized user views, working sets, and other user-definable features. Each time BrightStor Resource Manager is started, it prompts for the user ID and password, and loads the appropriate set of definitions.

If the database is stored in a central location and a link is present to the Administrator utility but you continue to work without creating new users, BrightStor Resource Manager does not prompt you for a user ID when you log in and continues to save your definitions in the **TD.MDB** file. If, at a later time, you want to set up individual user accounts, you must manually copy the **TD.MDB** file into your new **TD_xxx.MDB** file to preserve your existing definitions.

Use the Administrator utility to customize user definitions:

✓ Select the **BrightStor Resource Manager Administrator** utility from the BrightStor Resource Manager program group.

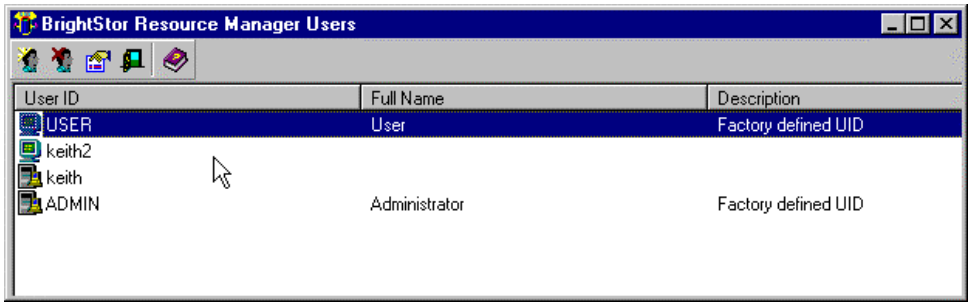

The utility displays a table showing the IDs, full names, and descriptions of the defined users. The following functions are available from icons at the top of the table or from the right mouse button pop-up menu:

- Add a new user
- Delete user
- Rename user
- Edit user properties

To create a new user:

1. Click the **Add a new user** icon or select **New User** from the right mouse button pop-up menu.

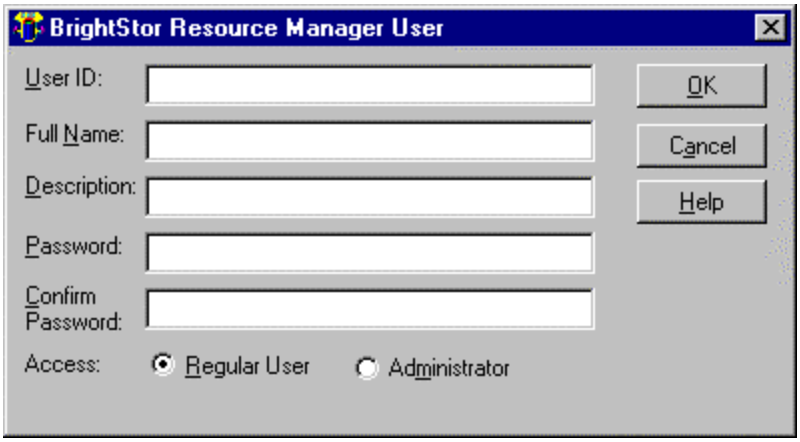

2. When the New User Properties dialog appears, enter user ID, optional full name, and optional description of the user in the appropriate fields.
3. Enter a password for the new user and type in the password confirmation (optional).
4. Click the appropriate button to specify whether the user is an administrator or regular user, and click **OK**.

The new user is added to the table of users.

To delete a user:

✓ Select the user you want to delete and click the **Delete the User** icon.

To rename a user (z/OS only):

1. Select the user you want to rename and click the **Rename the User** icon.
2. When the Rename User dialog appears, enter the new user name in the **Change to** field and click **OK**.

To edit a user definition:

1. Select the user you want to edit and click the **Edit user properties** icon.
2. When the Edit User Properties dialog appears, you can edit all the fields except the User ID. Click **OK** to save the changes.

To exit the BrightStor Resource Manager Administrator utility:

✓ Click the **Exit** icon.

Chapter

4

Navigating the Windows Interface

This chapter explains how to use BrightStor Resource Manager to view objects and manipulate object tables.

The BrightStor Resource Manager Windows Client provides a convenient graphic user interface to all functions of the BrightStor Resource Manager system. The user interface consists of the following main components:

- **Toolbar** - the most frequently used functions and a menu of general system functions and operations
- **Object tree** - a hierarchical view of all BrightStor Resource Manager objects
- **Hosts list** - all the hosts with which the client communicates
- **Consoles window** - an optional list of user-defined consoles that present at a glance information collected from several hosts
- **Working sets window** - an optional collection of frequently used objects
- **Active Objects window** - all active objects (open tables)
- **Scheduler window** - access to the BrightStor Resource Manager Scheduler
- **Agents list** - all the agents defined for the client
- **BrightStor Resource Manager tables** - the actual objects you are viewing and manipulating
- **Wizards** - to help you define trend logs, automation scripts, audit scripts, and system scripts (z/OS only)

Each of these components is described separately later in this guide.

To start the BrightStor Resource Manager client click the BrightStor Resource Manager client icon in the BrightStor Resource Manager group.

The first component that appears is the Toolbar, which provides quick, single-click access to the main menu functions.

Agents

The BrightStor Resource Manager system comprises various types of clients and servers. Both View 3270 and Windows Clients communicate with servers of the z/OS type. In addition, Windows Clients communicate with Windows servers to obtain network data.

Server types are represented in the Windows Client by BrightStor Resource Manager agents. The z/OS agent collects data from all the z/OS hosts to which the Windows Client is connected. The BrightStor Resource Manager agent collects data directly from the BrightStor Resource Manager database. The BrightStor Resource Manager agent fetches data exported by BrightStor Resource Manager to a shared directory. The Windows Client communicates with all servers of a certain type through the same agent. For example, the z/OS agent is used to communicate with all z/OS hosts.

BrightStor Resource Manager agents are installed by default together with the system. You cannot add or delete agents, but you can activate and deactivate them. You can perform these operations from the agent list.

Agent List

The agent list is invoked from the main menu.

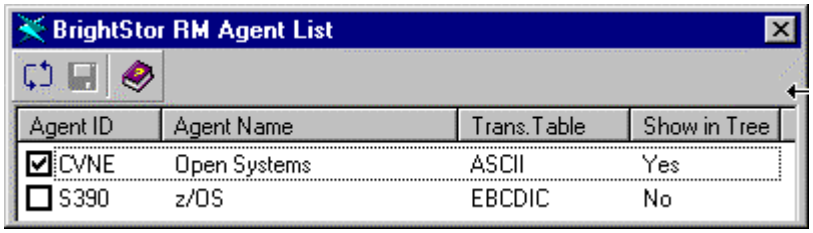

The following agent properties appear in the agent list:

Field	Description
Agent ID	Uniquely identifies the agent. Can be either CVNE or S390
Agent Name	Can be either Open Systems or z/OS.
Translate Table	Can be EBCDIC or ASCII. (z/OS only)
Show in tree	Can be **Yes** (show the source objects belonging to the agent in the object tree) or **No** (do not show), depending on whether the agent box in the Agent ID field is checked or not.

z/OS agents are active by default (the box next to the Agent ID is checked). You can deactivate an agent by unchecking its box. The client receives data only from servers that belong to active agents. If you deactivate the agent then data collection stops from all the servers of the agent's type.

Agent Properties

Agents also appear elsewhere in the user interface, in the object tree and in the host definition forms, where they interact with other system functions.

Objects in the BrightStor Resource Manager object tree appear under the agents that collect their data. The object tree shows only objects belonging to active agents (whose box on the agent list is checked). If you uncheck the box next to an agent, its objects are removed from the object tree.

To view agent properties select an agent in the object tree and click the right mouse button, then choose **Properties**.

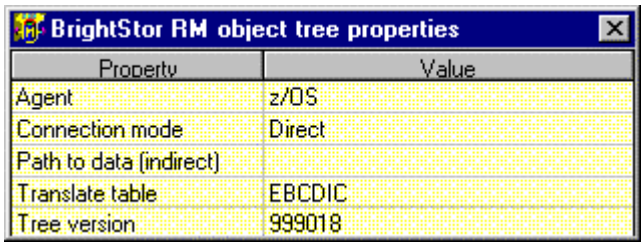

The agent properties screen contains the following information:

Field	Description
Agent	Agent name
Connection mode	The type of connection to the servers. Can be **Direct** or **Indirect**.
Path to data	The path to the shared directory (for indirect connections only).
Translate table	Can be EBCDIC or ASCII
Tree version	Version of the object tree currently displayed under the selected agent. To change the version of the object tree, see Hosts Menu.

Agents and Hosts

The Windows Client displays data collected from a single Open Systems host (Application Server) and any number of z/OS hosts (servers) simultaneously. One of the properties of each host is the name of the agent that collects data from it.

Every host in the BrightStor Resource Manager host list must belong to one of the agents installed with the system. A host can belong to only one agent, but many hosts can belong to the same agent.

Hosts that belong to different agents have different properties and their data is collected and transferred to the Windows Client by different methods.

Host List

A BrightStor CA-Vantage Storage Resource Manager host is a server that furnishes data to a Windows Client. Each host is identified by the agent that collects the data, and in the case of z/OS agents, by an IP address and a port number. To log into hosts belonging to z/OS agents, a user ID and password are also needed, but these can be supplied at login.

The Host List window lists the hosts that have been defined on the Windows Client.

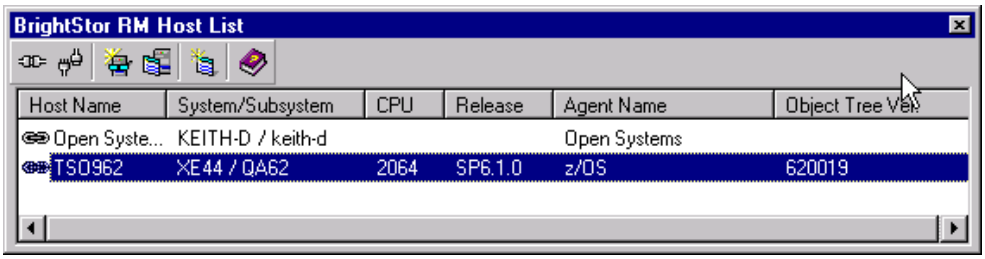

You can perform the following operations with hosts:

- Add a new host using the wizard (z/OS only)
- Add a new host using the viewer form (z/OS only)
- Connect/disconnect from host
- Select functions from the menu that appears when you click the right mouse button

Adding Hosts

You can add new z/OS hosts by selecting **New host** from the right mouse button menu, and then selecting one of the two options: **By Wizard** or **By Viewer**.

Adding Host by Wizard

The Host definition wizard guides you step-by-step through the procedure of creating a new z/OS host.

1. Enter a name for the new host.
2. Select an agent running on the host from the **Agent Name** drop-down list. If you selected an agent with a direct connection mode (z/OS), click **Next**. If you selected an agent with an indirect connection mode (Open Systems), click **Finish**.
3. If you selected an agent with a direct connection mode, enter the IP address and port number of the mainframe to which you want to connect, then click **Next**.
4. Optionally, enter system and subsystem IDs for the mainframe to which you want to connect, then click **Next**.
5. Optionally, enter your user ID and password on the mainframe to which you want to connect. Re-enter the password to confirm. Click **Next**.

   If you leave the **User ID** and **Password** fields empty, each time you connect to the host, BrightStor Resource Manager displays a login dialog asking for your user ID and password.

   If you enter a user ID but no password during host definition, each time you connect to the host, BrightStor Resource Manager displays a login dialog asking for the password.

   Note: Both User ID and password are encrypted and sent to the host in this form.
6. Optionally, in the **Host Trace** field, enter the value EWS10T to create a trace on the host of all the messages it sends to and receives from the program.
7. Optionally, check the **Communication logging** box to save messages relating to the host-client communication in a log file on your workstation. Click **Next**.
8. Optionally, click **Test** to verify the IP address and port number you specified.
9. Click **Finish** to complete the definition.

This ends the action of the wizard. The created host is added to the Host list. To review the definition and make changes if necessary, select **Host definition** from the right mouse button menu. The Host definition form shows all your parameter settings.

To review the definition and make changes if necessary, select **Host definition** from the right mouse button menu. The host definition form shows all your parameter settings.

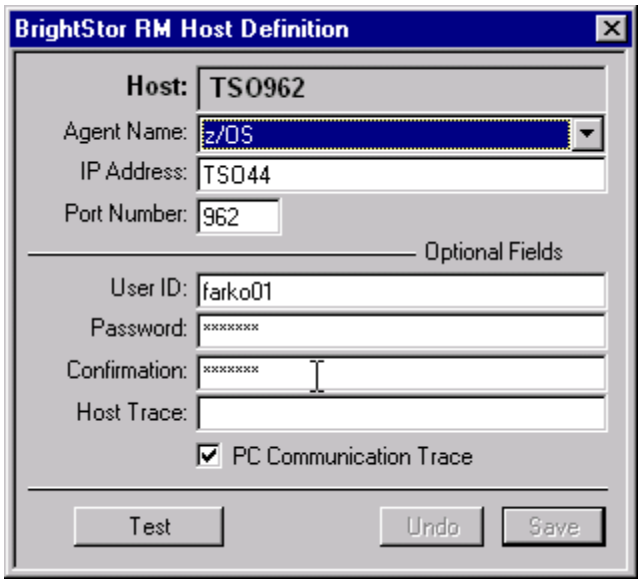

You can modify any of the parameters, except the host name. Click **Save** to apply the host definition changes you make on this form.

Adding Host by Viewer

To define a new z/OS host or to modify an existing one:

1. Select **New Host** from the right mouse button menu of the Host List and select **By Viewer.**

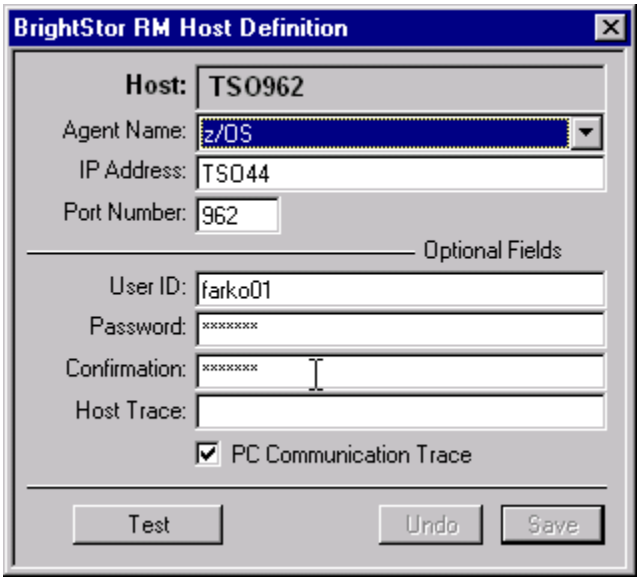

2. Enter a name for the new host in the **Host** field.
3. Select an agent running on the host from the **Agent Name** drop-down list.
4. If you selected an agent with a direct connection mode, enter the IP address and port number of the mainframe to which you want to connect in the appropriate fields.

   The remaining fields are optional. If you leave the **User ID** and **Password** fields empty, each time you connect to the host, BrightStor Resource Manager displays a login dialog asking for your user ID and password. If you enter a user ID but no password during host definition, each time you connect to the host, BrightStor Resource Manager displays a login dialog asking for the password.
5. Optionally, in the **Host Trace** field, enter the value EWS10T to create a trace on the host of all the messages it sends to and receives from the program.
6. Optionally, check the **PC Communication Trace** box to save messages relating to the host-client communication in a log file on your workstation.
7. Optionally, click **Test** to verify the IP address and port number you specified.
8. Click **Save** to save the definition or **Undo** to clear the form.

Note: Once you enter data in any of the fields, you cannot exit the dialog until you click either **Undo** or **Save**.

Host Definition

The host definition form contains the fields shown in the following table. The table also indicates which fields are required and which are optional for agents that maintain a direct connection to the host (z/OS) and an indirect connection (Open Systems).

Field	Description	Required
Host	Name of the host you are about to define.	Both
Agent Name	Name of the agent that collects data from the host.	Both
IP Address	IP address or server name of the z/OS host to which you are connecting.	Direct only
Port Number	Port number of the z/OS host to which you are connecting.	Direct only
User ID	Your user ID. It is sent to the z/OS host in encrypted form.	Neither
Password	Your password. It is sent to the z/OS host in encrypted form.	Neither
Confirmation	Confirmation of your password.	Neither
Host Trace	The value EWS10T creates a trace on the z/OS host of all the messages exchanged between the host and the Windows Client. Use only if explicitly instructed by technical support staff.	Neither
PC Communication Trace	If checked, the system saves the messages relating to host-client communication in a log file. The file has the same name as the host and a .DAT suffix.	Neither

Connecting to the Host

To connect to a host double-click the desired host or select the host and click the **Connect** icon in the Host List toolbar.

The host icon changes to indicate that the client is trying to connect to the host. When the connection is established, the icon changes again to show a connected host.

Note: If only one host is defined, it is always selected, and the client tries to connect automatically at start-up. If more than one host is defined, connection at start-up is controlled by settings in the Hosts panel of the Activation dialog, under the **Tools** menu.

Hosts can have the following statuses:

- Connected
- Not connected
- Connecting

Each status is identified by a characteristic icon.

Operations performed on BrightStor Resource Manager objects listed in the object tree are always with reference to one or more hosts. To open any object you must first specify the host from which you want the data to be retrieved. You do so by selecting a host from the list. If more than one host is selected when you request a view in the object tree, an object is opened for each selected host. If a host is selected but not connected when you issue a request for a view, the Windows Client tries to connect to the host.

You can select a host by clicking it. You can also use the right mouse button menu to select several hosts with one command:

✓ On the right mouse button menu, choose **Select Hosts**, then **All** (to select all the hosts in the list), **Connected** (to select all connected hosts), or **Disconnected** (to select all disconnected hosts).

If only one host is defined, you need not select it. The system assumes that it is always selected.

Hosts Menu

Clicking the right mouse button anywhere on the Host List brings up the following menu.

To connect or disconnect one or more hosts:

✓ Select the hosts you want to connect/disconnect, then click the right mouse button and choose **Connect** or **Disconnect**.

To select more than one host at the same time:

✓ From the right mouse button, choose **Select Hosts**, then **All** (to select all defined hosts), **Connected** (to select all currently connected hosts), or **Disconnected** (to select all disconnected hosts).

To define a new z/OS host:

✓ From the right mouse button, choose **New host**, then **By Wizard** or **By Viewer**.

To define a new host by modifying an existing one:

1. Select the host you want to clone, then click the right mouse button and choose **Save As**.

   When the Host Definition dialog appears, all the fields contain the values of the selected host but the **Host Name** field is empty.

2. Enter a new name for the cloned host, make whatever changes are necessary to the other fields, and click **Save**.

To remove a host:

✓ Select the hosts you want to remove, then click the right mouse button and choose **Remove host**.

To modify an existing host definition:

✓ Select the hosts you want to modify, then click the right mouse button and choose **Host definition**. Modify the definition as needed.

To change your user ID, password, or both:

1. Click the right mouse button and choose **User IDs and Passwords**.
2. On the Define User IDs dialog that appears, you can enter a new user ID and new password in the appropriate fields. After you click **OK**, the system updates the user IDs and passwords on all the hosts that use the same user ID.

To change the object tree version displayed on your client machine to correspond to the BrightStor Resource Manager release running on the selected z/OS host:

1. Select the desired host on the host list.

2. From the right mouse menu of the host list, select **Object Tree Services**, then **Display Object Tree**.

   If the object tree version corresponding to the BrightStor Resource Manager release running on the selected host is available on the client machine, it is immediately displayed, replacing the current object tree version running on the client. If the object tree version is not available on the client machine, it is first downloaded from the host, then displayed. A progress bar shows the status of the download.

To change the object tree so that it shows only objects to which the logged-in user has access on the selected z/OS host:

1. Select the desired host in the host list.
2. From the right mouse menu of the host list, select **Object Tree Services**, then **Display Available Objects**.

To download the object tree corresponding to the BrightStor Resource Manager release running on the selected host even if it is available on the client machine:

1. Select the desired host on the host list.
2. From the right mouse menu of the host list, select **Object Tree Services**, then **Reload Object Tree**.

Note: Use this option to restart the download if the object tree was not displayed properly because of a communication problem.

If the object tree window is open, the refreshed version is displayed as soon as the download completes.

To refresh the host list to show the latest additions and deletions:

✓ Click the right mouse button and choose **Refresh Host List**.

Object Tree

The expandable/collapsible object tree lists all the **source objects** known to BrightStor Resource Manager. A source object contains all the fields collected by the host for that object. All fields or any subset of them can participate in user views based on that object. Source objects appear in the object tree in green. User views appear in blue.

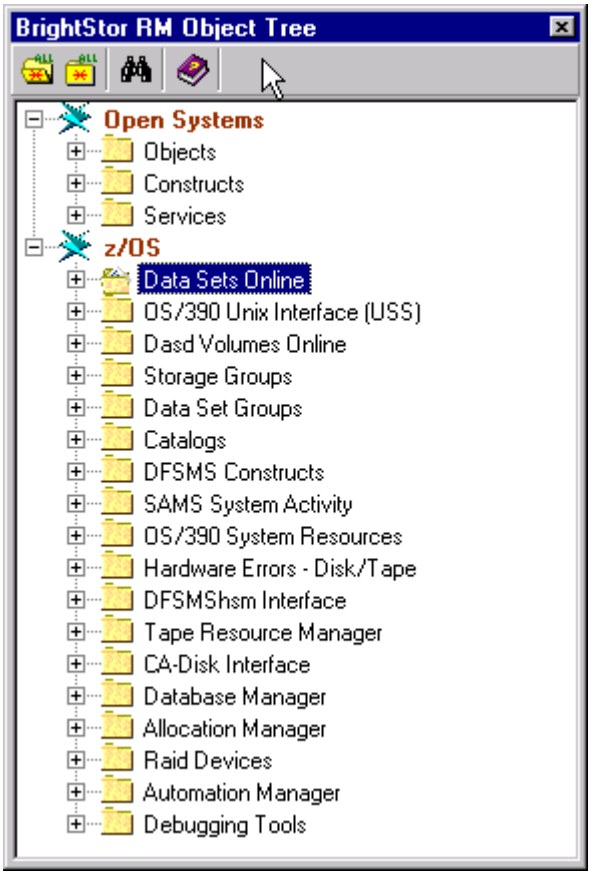

Source objects are assembled on the object tree into folders that represent categories of objects. To expand a folder:

✓ Double-click the folder icon.

The folder expands to show a hierarchy of objects. If the hierarchy contains several levels, you proceed to the lower levels by double-clicking the folders that contain the lower levels.

You cannot modify or delete a source object. But you can create your own user views based on a source object.

User views reside under the source objects to which they belong. A user view consists of a subset of object attributes. Any number of user views can be derived from each source object. You can delete and modify user views as needed.

The BrightStor Resource Manager installation is delivered with a selection of pre-defined user views. With each new release of BrightStor Resource Manager, a new set of user views is supplied on the installation media, but not integrated into the object tree. You can use the Import function to bring the new user views into the object tree.

You can modify and save these views under different names to create your own views. All user views appear in the object tree under the source objects from which they are derived.

In every respect other than the way in which they are created and modified, source objects and user views behave identically.

To open a source object or a user view:

- ✓ On the Host List, select the host or hosts from which you want to retrieve the data, then click the desired object or view.

  If you select more than one host, BrightStor Resource Manager opens a table with the desired object or view for each selected host. If a host is selected but not connected, the system attempts to make the connection.

  If only one host is defined in your Host List, you need not select it before clicking the desired object or view; the host is selected by default.

To collapse a folder:

- ✓ Double-click the open folder icon.

Object Tree Menu

Clicking the right mouse button anywhere on the object tree brings up the right mouse button menu. The menu varies according to the current selection in the object tree.

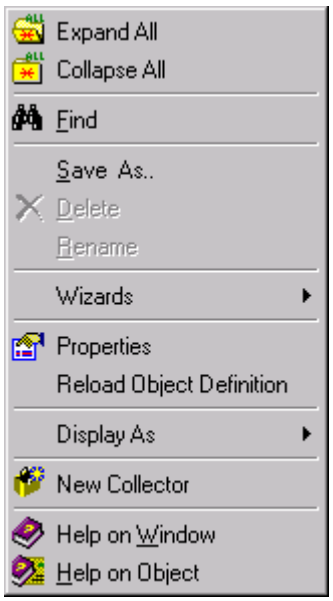

To expand or collapse the hierarchy of views below an object:

✓ Select the object or view you want to expand or collapse, then choose **Expand** or **Collapse** from the right mouse button menu.

Alternatively, select **Expand All** or **Collapse All** to fully expand or collapse the tree.

To find an object or a view in the tree:

1. Choose **Find** from the right mouse button menu.
2. When the Find dialog appears, enter the name of the object or view in the **Find What** field.
3. Choose **Down** or **Up** from the **Direction** drop-down list to indicate the direction of the search. Optionally, check the **Find whole word only** and **Match case** boxes to restrict the search.
4. Click **Find Next**.

   The system places the cursor on the next occurrence of the search string in the object tree.

To save an existing view under a new name:

1. Select the desired view and choose **Save As** from the right mouse button menu.
2. When the Save dialog appears, enter a new name for the selected view in the **Name** field and click **Save**.

To delete a user-defined view:

✓ Select the view you want to delete, then choose **Delete** from the right mouse button menu.

To rename a user-defined view:

1. Select the view you want to rename, then choose **Rename** from the right mouse button menu.
2. Specify a new name for the view and press **Enter**.

To start a z/OS object logging or automation wizard:

1. Select the desired view and choose **Wizards** from the right mouse button menu.
2. When the secondary menu appears, select **Object Logging** (see the section Defining Log Operations).

To control the color coding of individual fields in a view or to change the scale base of columns containing numeric data:

✓ Select the view, then choose **Properties** from the right mouse button menu. Refer to Object Properties for details.

To reload the definition of a selected view from the host:

✓ Select the view, then choose **Reload Object Definition** from the right mouse button menu.

This function is useful if an object definition is changed on the host.

To display a user view that is not a table (for example, graphs and reports), *as a table*:

✓ Select a user view, then choose **Display As** from the right mouse button menu. From the secondary menu, select **Display as Table** (real data) or **Display as Demo Mode Table** (view data saved on the client using the **Save Data for Demo** function (see Save Data for Demo)). Use this function to make changes to a non-tabular user view.

For z/OS only, to define a new trend collector or new trend report, or to modify the properties of existing ones:

✓ Select **New Collector**, **New Trend Report**, **Collector Properties**, or **Trend Report Properties** from the right mouse button menu. Refer to Trend Reports (z/OS only).

Accessing an Object

To open one of the BrightStor Resource Manager objects:

1. Select one or more hosts from which you want to obtain the data.
2. On the object tree, click the object you want to open. You can expand and collapse the object tree to bring the desired object into view.

For each selected host, the system opens a table showing the object data on that host. If you selected hosts that were not connected, the system tries to connect to those hosts as well.

Tables

After you select an object, and after the table used to display information about the selected object appears, a menu becomes available for the table. The following functions are available for all objects and views. Additional functions, accessible from individual views only, are described under those views. Some of the View Menu functions are only available for z/OS objects and some menu items may vary with the objects.

Title Menu

- Open (z/OS)
- Open Zoom (Open Systems)
- Print
- Print Preview
- Print Definition
- Export
- Save/Save As
- Save Data for Demo

View Menu

- Sort
- Filter
- View Definition
- Mode (z/OS)
- Monitor (z/OS)
- Graph
- Totals (z/OS)
- Statistics (z/OS)
- Execute
- Refresh
- Rebuild Table

Edit Menu

- Select/Deselect All Lines

Actions Menu

In some views you can take actions on selected records (rows). The actions available for the various objects and views differ. The online help contains information about the specific actions available for each view. For a general description of actions, refer to Actions Menu.

In addition to the menu options, there are several functions accessible only from buttons in the toolbar. These are:

- Display Detailed Line
- Copy to Clipboard
- Show line, used to go directly to a specified line in the table
- Adjust column width, used to adjust each column to the width of the current data in it

The Table Window

BrightStor Resource Manager Objects are presented in configurable tables. You can change both display characteristics (such as the width and number of the table columns, the number of rows displayed, and so on) and characteristics that apply to the object displayed in the table, such as sort and filter criteria, for example. Together they determine the table's appearance. Both display and object characteristics can be saved as part of a user-defined view.

BrightStor Resource Manager data is organized in table windows similar to the one displayed below.

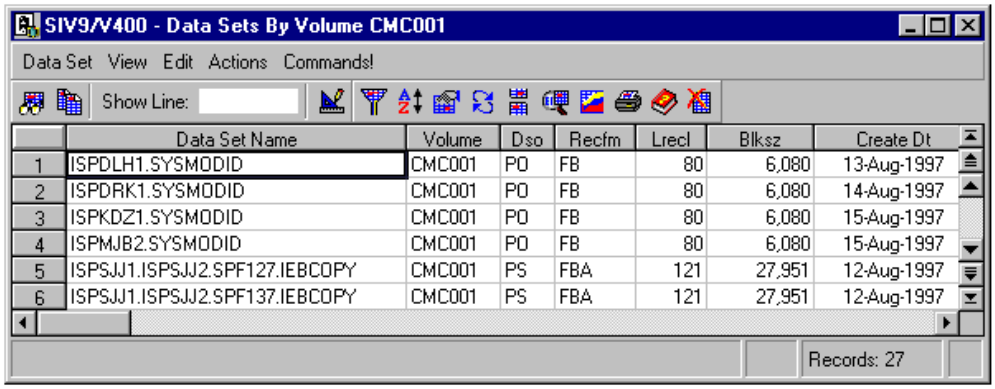

In the table window you can find the following elements:

- The window title (or caption) identifies the host system and subsystem from which data comes and describes the source object from which the data is derived.
- The left part of the status line at the bottom of the table shows the name and description of the source object from which the data is derived. For user-defined views, the status line also shows the view name and description. Initials in the right part of the status line indicate the mode that was used to collect the data (realtime, last interval, snapshot copy, disk checkpoint, log, or capture). If the data is last interval data, placing the cursor on the LI initials displays the time of the last sampling of the object. The number of records found is displayed at the far right of the status line.
- Quick-access buttons (icons) in the toolbar provide easy, one-click access to the most common functions.
- Rows and columns of information about storage management objects make up the body of the table. Each row corresponds to an object instance and is numbered in ascending order along the left edge of the window. Each column displays information about a particular attribute of the object.
- **Show Line** box can be used to go directly to a particular row in the table.
- The first column of the table, generally containing the name of the records that appear in the table, remains in view and does not scroll away while you scroll the table columns horizontally.

To exit the object table:

✓ Click the **Close** button in the table's toolbar or select **Close** from the table's first menu.

Resizing Table Columns

You can resize columns interactively:

1. After table data is displayed, move the mouse pointer into the column's heading line and place it on the line between the headings of two columns.

   The shape of the pointer changes to a double arrow pointing left and right.

2. Drag the mouse to the left or to the right to decrease or increase the width of the column to the left of the cursor.

   - or -

✓ Click the **Adjust column width to data width** icon in the toolbar to automatically adjust the width of all the columns in the table to the width of the data.

You can move, resize, and reposition the tables that BrightStor Resource Manager displays on the screen. Next time you open the same window, it appears with the size and in the position it was when you closed it.

You can globally resize table columns from the **Activation** function of the main BrightStor Resource Manager menu.

1. Select **Activation** from the **Tools** menu.
2. Check the Automatic Table Column Adjustment box on the General tab.

Next time you start BrightStor Resource Manager, every column in every table is automatically adjusted to the width of the data when the table is opened. When this option is selected, the Adjust Column Width button is not shown on the Toolbar.

Display Detailed Line Button

In the upper left corner of the table, just above the row number, is the **Display Detailed Line** (Eyeglass) button used to display all the information of a selected table row in a single window. You can use this button to view all the columns of a wide table at once, eliminating the need to scroll laterally through the columns of the table.

Use the **Prev** and **Next** buttons of the full detail window to view the previous and next record in the table.

View Definition

Use this function to create user-defined views based on source objects. The dialog consists of the following tabs: **View**, **Data Request**, **Fields**, **Schedules**, **Hosts**, and **HTML Properties**. Click **OK** after you make your selections in all the tabs (you can click **OK** from any tab).

View Tab

Use this tab to specify the type of the view. The options are:

- Table window: the standard BrightStor Resource Manager window that displays data in tabular form.
- Graph window: a graphic presentation of the data. Use the Graph function to define this type of view.
- Light trend report: a graphic presentation of trend data, which can be selected only when in a Log window. Use the Trend Definition function to define this type of view.
- Printout: data formatted for print. Use the Print Definition function to define this type of view.
- HTML file: object data exported to HTML format and saved under a tree of HTML user views. At the root of the tree is the file **DEFAULT.HTM**, in the BrightStor Resource Manager directory. You can access the HTML user views from any client machine through the Internet using a standard browser.

  Notes:

  - To make the HTML file available to remote clients, you must run either IIS or Personal Web Server on the BrightStor Resource Manager machine that hosts the HTML file and define a virtual directory on the server to point to the BrightStor Resource Manager directory. Refer to your Windows documentation for details.
  - To view the HTML file, you must set up the HTML user view browsing ability on the client machine.
  - The HTML files generated by BrightStor Resource Manager use Active X controls. You need Microsoft Explorer to view the HTML pages exported by BrightStor Resource Manager. If you use Netscape or another browser, you must install the ActiveX plug-in to view the data.

If you selected more than one host on the Hosts tab, you must specify here how to display the data of the various hosts. The available options are:

- Open a separate table for each selected host. This is the default option.

- Show the data from all hosts in one table, without consolidating identical records. In this case, a new column, HOST, is added to the table to specify the host from which each record has been retrieved.
- Show the data from all hosts in one table, and consolidate identical records. In this case, identical records are shown only once, and a tilde (~) sign is entered in the HOST column for consolidated records.

Note: If you choose to consolidate identical records, make sure to save the user view under a new name using the **Save As** function.

Data Request Tab

Use this tab to specify the method used to request data from the host and the number of records requested.

- Specify whether the first data request is issued manually, after clicking the **Execute** button, or automatically, when a view is opened. The manual method lets you define sort and filter functions before issuing the request, which can result in a reduction of the amount of data transferred.
- For large tables, you can choose to request all the records from the host with the first request, or to request a specified number of records only.

Selections made in the **Data Request** tab take effect the next time you start BrightStor Resource Manager, without the need to save the view.

Fields Tab

Use this tab to select which fields are displayed in a view.

- ✓ Use the left and right arrows to move each field between the **Available Fields** and the **Selected Fields** list. You can use the double arrows to move all the fields into one or the other list.

Your selections take effect immediately.

Schedules Tab

Use this tab to define schedules. When you select the Schedules tab for the first time (and no schedule has yet been defined), the panel appears empty. To define a schedule:

1. Click the **Add** button.

2. When the New scheduler dialog appears, define the desired schedule by making the appropriate selections from the **Every** (number of minutes, hours, days, weeks, months), **On** (Monday, Tuesday, etc.), and **At** (time) drop-down lists. A typical schedule may have the following format: Every Week on Friday at 17:00.
3. Click **OK**.

   When the Schedules tab is redisplayed, the defined schedule is listed in the Schedules window, with its complete description, for example, Every 8 Weeks on Tuesday at 08:00.

To add another schedule:

✓ Click the **Add** button once more and repeat the operation. You can add as many schedules as needed.

To delete a schedule:

✓ Select the schedule and click the **Delete** button.

To modify the definition of a schedule:

✓ Select the schedule and click the **View** button. When the Scheduler dialog appears, make the desired changes and click **OK**.

After you complete all the definitions and click **OK**, the system displays the following message:

```
This change will be activated the next time the User View is opened. Don't forget
to save your User View to keep this change.
```

✓ Use the **Save** or **Save As** functions from the object table to save the user view.

Note: Make sure the Scheduler is active. The schedule you have defined is executed only if the Scheduler is running at the time the scheduled action is due for execution (see Scheduler).

Hosts Tab (z/OS)

Use this tab to specify the hosts from which data is retrieved for scheduled operations.

Note: If you select more than one host, you must choose on the View tab the method to be used to display the data of the various hosts.

1. Specify the host or hosts from which to retrieve data for scheduled operations.

   Clicking the **All hosts** button selects all hosts defined in the system.

Clicking the **Selected hosts** button selects all hosts currently selected in the Host List.

Clicking the **Connected hosts** button selects all hosts currently connected.

Clicking the **Specific hosts** lets you check the boxes of the desired hosts in the list below.

2. After completing the selection process, you can check the **Show selected hosts only** box to display the list of selected hosts only.

HTML Properties Tab

Note: The HTML files generated by BrightStor Resource Manager use Active X controls. You need Microsoft Explorer to view the HTML pages exported by BrightStor Resource Manager. If you use Netscape or another browser, you must install the Active X plug-in to view the data.

The HTML Properties tab contains two subtabs:

- Use the **Categories** tab to define a subset of BrightStor Resource Manager object data exported in HTML format.
- Use the **Format** tab to add an optional header, title, comment, and footer to the exported data.

The tab is accessible only when the view type is **HTML file**.

Categories

An HTML category is a hierarchy of directories that contain a subset of BrightStor Resource Manager object data exported in HTML format. You define the root of the directory for each category you create.

The data of all BrightStor Resource Manager objects and views is exported in HTML format to a default hierarchy of files, with its root in the file DEFAULT.HTM, in the BrightStor Resource Manager directory. *In addition,* the data of objects and views associated with user-defined HTML categories is saved under a hierarchy of files, as defined by each HTML category.

Use this tab to create new categories of objects exported in HTML format and to associate the current view with one or more HTML categories. A category can be stored in its own directory.

A view can belong to several categories. A category, in turn, can contain several views of the same object or of different objects. When BrightStor Resource Manager executes the export schedule, objects belonging to active categories are exported to their respective directories, in addition to the default export to the BrightStor Resource Manager directory. Categories can be activated and deactivated as needed. This provides a convenient way to start and stop the export of multiple objects and views, and helps control user access to exported view data through IIS security.

To create a new category:

1. Click **Add**.
2. When the Add Category dialog appears, enter the category name, a description, and navigate to the path of the directory where you want the exported HTML data to reside. Click **Add**.

   The new category is added to the list in the HTML Categories tab, with a checkbox next to its name.

Note: Each category you add becomes available to all BrightStor Resource Manager objects and can be selected from the View Definition dialog of every object.

To edit an existing category:

✓ Select the desired category and click **Edit**. The Edit Category and Add Category dialogs are identical.

To delete category:

✓ Select the desired category and click **Delete**.

To export view data to the directory specified for a given category:

✓ Check the box next to that category. If you check several categories, view data is exported to each directory specified in the checked categories.

Format

The Format tab lets you add a header (top line), footer (bottom line) and user note (text that appears above the bottom line) to the HTML export data. All three elements contain free user text. You can also determine the alignment of the table title.

1. Enter the text of the desired element in its appropriate box.
2. Select **Left**, **Center**, or **Right** alignment for each element by clicking the corresponding button.
3. Click the **Color** button and set the desired color, where appropriate.

Mode (z/OS)

BrightStor Resource Manager collects data in the following timing modes:

- **Real Time** - the information displayed in the table is collected immediately after clicking **Execute**.
- **Last Time Interval** - the information displayed in the table was collected during the last regular scanning cycle on the host. The frequency of scanning for different objects is set in the corresponding parameters in VKGPARMS.
- **Snapshot Copy** - the information displayed in the table is a snapshot copy of a continuously updated object that changes dynamically as individual records within the object are being updated. When a snapshot copy is requested, BrightStor Resource Manager locks the object for as long as it takes to filter and copy the data. After the data is copied, BrightStor Resource Manager releases the object for further updates. The snapshot copy is then available for viewing while the original object is being updated.
- **Disk Checkpoint** - the information displayed in the table is taken from a checkpoint version written to disk at scheduled intervals. The disk checkpoint version is created by a script. Sample scripts are distributed in the VANTAGE.SAMPSCRI.PDS. Refer to the Object Definitions table under Debugging Tools to identify the script used to create the disk checkpoint for the object of your choice. Disk Checkpoint is available to enhance system performance in accessing selected objects such as CA-1 TMC volumes.
- **Log** - the information displayed in the table is taken from the object's log file. This information is available only if log data has previously been saved for the given object, using the **Object Trend Logging** in the **z/OS** menu (see Defining System Scripts for information on how to log data for a list of selected objects).
- **Capture** - the information displayed in the table is taken from the object's capture file. This information is available only if capture data has previously been saved using the **Export** function (see Export).

Note: When the capture or log data is displayed, the first columns of the object table contain date and time stamps that show when the data was collected and the system ID. You can filter the log and capture data according to specific dates and times.

To select the collection mode for the table data:

1. Select **Mode,** then one of the collection mode options from the object table's **View** menu.

If you selected **Log** or **Capture**, the Log/Capture Collection Period dialog appears. It lets you select the range of dates for the log or capture data.

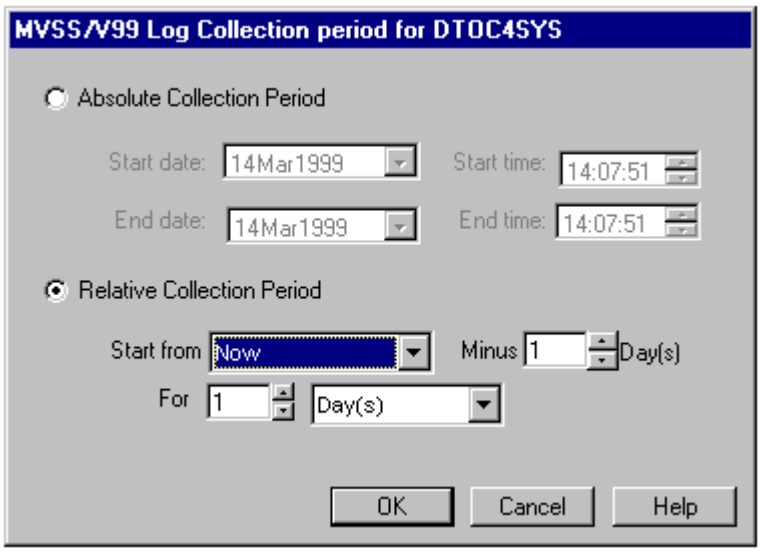

2. To define an absolute collection period, enter the desired time interval in the **Start date** and **End date** text boxes. Use **Start time** and **End time** boxes to enter the hour of the start and end of the time interval.

   When you click the down arrow of the drop-down box for either **Start date** or **End date**, a calendar appears to help you specify an absolute date:

   a. Click on the year to obtain a drop-down list box you can use to specify a year.

   b. Click on the month to obtain a drop-down list box you can use to specify a month.

   c. Click the left or right arrows to show the calendar for the previous or next month.

   d. Click on the desired day to select it.

   To define a relative collection period, enter a start date and a duration, both defined in relation to some date. Define the start time as a given time (now, midnight today, the start of the week) minus a number of days, weeks, and so on. Define the duration by indicating a number of minutes, hours, days, weeks, months, or years that you want collection to continue.

3. Click **OK**.

Note: Data Collection mode settings are saved as part of user-defined views.

Input List (z/OS)

The Input List dialog lets you specify the data sets from which object tables are created. If the data sets used to build the tables are specified in VKGPARMS, the Input List dialog overrides the VKGPARMS definition for the current session.

Data set names in Input Lists can be specified symbolically, so they can contain substitution variables as well as relative or absolute GDG numbers.

Note: Because an input list is not executed in any specific object context, you cannot use Object Related variables. You must use System Related variables or System Symbols.

To edit the input list:

1. From the **View** menu of one of the CA-Disk Interface or DFSMShsm Interface tables, select **Input List**.

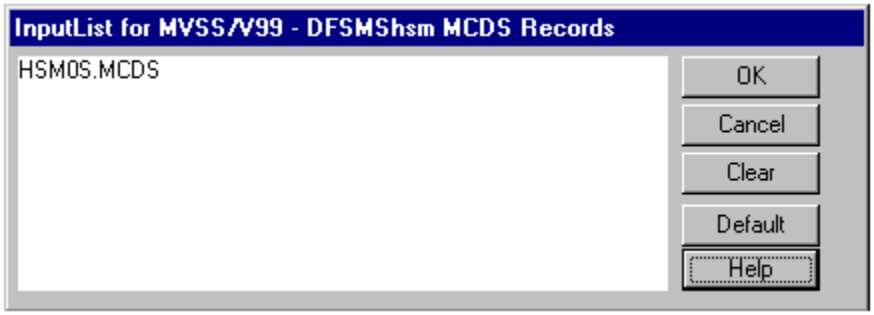

2. In the edit window that appears, type the name or names of the data sets from which you want to generate the given table.

   If VKGPARMS specifies default data set names to be used in creating this table (as in the case of CA-Disk tables, whose input files are listed under FILELIST or FILES1 through FILES15 in the System Parameters display for CA-Disk parameters), you can view the default data set names. To do so:

3. Click the **Default** button.
4. Review the list of data sets that appears in the window and edit it.
5. Click **OK** when done.

Filtering

Use the **Filter** function to narrow the list of objects displayed in the table.

To invoke the Filter dialog:

✓ Select **Filter** from the object table's **View** menu or click the **Filter** button.

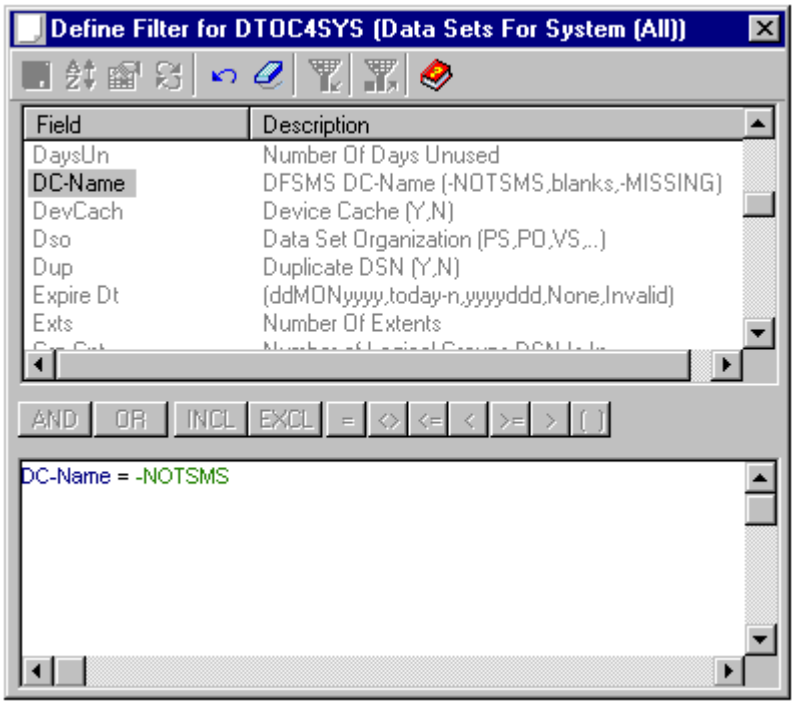

Object attributes are listed in the Filter window in alphabetical order by field name. You can click at the top of the **Description** column to sort entries by description name. You can reverse the sort order by clicking again at the top of either column.

You can build a filter from the attributes of the objects displayed in a window by combining them in Boolean expressions. In the case of volumes, for example, objects have such attributes as **Volume Name**, **Device Number**, **Device Type**, **Model Type**, and so on. Filter expressions can contain patterns created with wild card characters (see Wildcard Characters for more information). You can use single quotes to filter text strings with blanks or any values containing blanks. If blanks are a significant part of the substring for which you are searching within a text field, you can use double quotes rather than single quotes.

When filtering on date fields, you can compare against either an absolute date in Gregorian (ddMONyyyy) or Julian (yyyyddd) format, or one calculated by adding or subtracting nnn days from the current date (see Filtering on Date Fields for more information).

You can refine your filter by using the AND and OR logical operators to combine several expressions. You can also use parentheses to group sub-expressions.

The Filter dialog guides you in the process of defining the filter expression by enabling and disabling the appropriate fields and controls at every step.

You can enter a filter expression directly into the text box at the bottom of the Define Filter dialog, or use the typing aids available in the dialog. Whichever way you created the filter expression, you can edit it in the text box.

To define a filter expression using the typing aids:

1. Click the desired attribute (field) name.

   The attribute name is transferred to the filter text box at the bottom of the window, and the appropriate **operator** buttons are enabled.

2. Click the desired **operator** button to select a relational operator.

   The relational operator is transferred to the filter text box.

3. Type in the desired value or pattern in the filter text box.

   To incorporate spaces in the value, enclose them within single quotation marks.

4. Click the desired **operator** button to select a desired logical operator.

5. Repeat Steps 2 through 4 if necessary, to complete your filter expression.

To gradually reverse the steps of the filter expression definition:

✓ Click the **Undo** button.

To open the Sort dialog:

✓ Click the **Sort** icon.

To execute the filtering action:

1. Click the **Set and Exit** icon.

   If the filter passes the syntax check on the host, the dialog closes and the table from which you invoked the filter becomes active.

2. Click **Execute** to apply the filter to the table.

   - Alternatively-

✓ Click **Execute** in the Filter dialog. This function is the equivalent of clicking **Set and Exit** in the Filter dialog and **Execute** in the table.

In some cases, clicking **Execute** initiates a new data collection on the host. Refer to Execute to learn about the cases in which the **Execute** function initiates a new data collection.

To force a new data collection and then execute the filtering action:

✓ Click **Refresh**.

Note: If you also want to sort your data, do not click **Execute** at this point. First define your sort criteria by invoking the Sort dialog, then click **Execute** in the Sort dialog.

To save a filter definition as an external filter (see External Filters (z/OS)):

1. Create a filter using the standard procedure.
2. Click the **Save as external filter** icon on the toolbar.

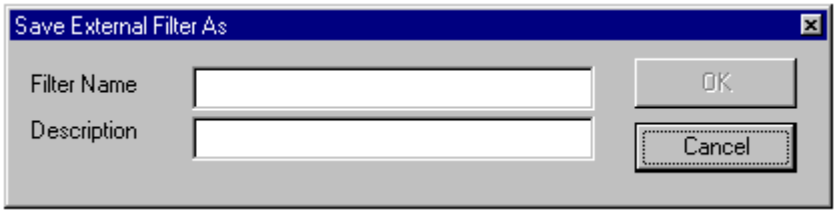

3. When the Save External Filter As dialog appears, enter a name and description for the filter and click **OK**.

To import an external filter from the common global repository on the host:

1. Click the **Insert external filter** icon on the toolbar.

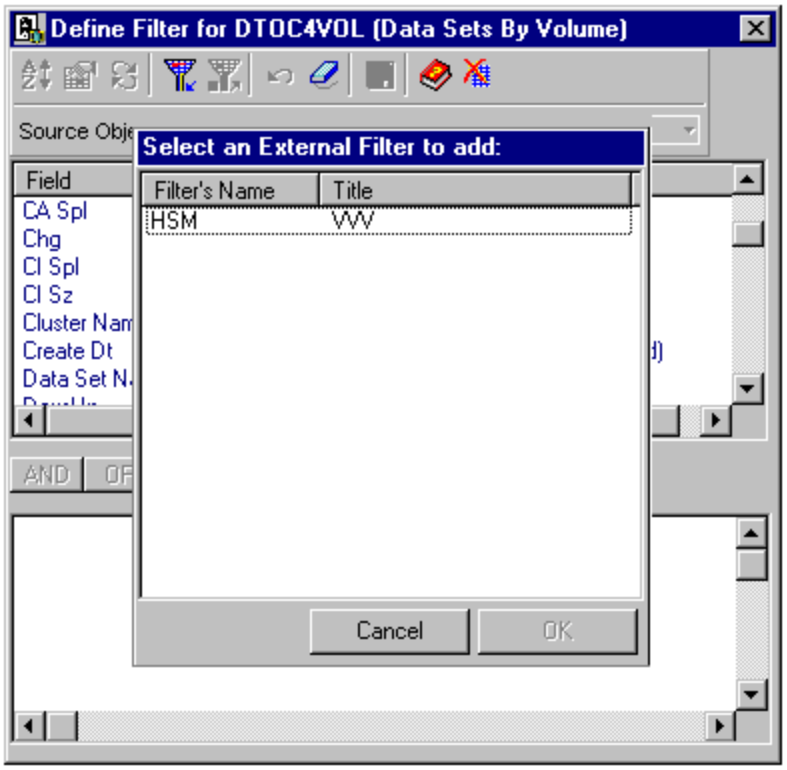

2. When the list of available external filters appears, select the desired filter and click **OK**.

   The selected filter definition is brought into the Define Filter window at the position of the cursor, where you can use it as is or modify it according to your needs.

External Filters (z/OS)

You can save filters in a common global repository on the **z/OS** host and make them available for reuse by any user.

You can add, edit, or delete filters from the External Filter dialog. To invoke the dialog:

✓ Select **External Filters** from the **z/OS** menu.

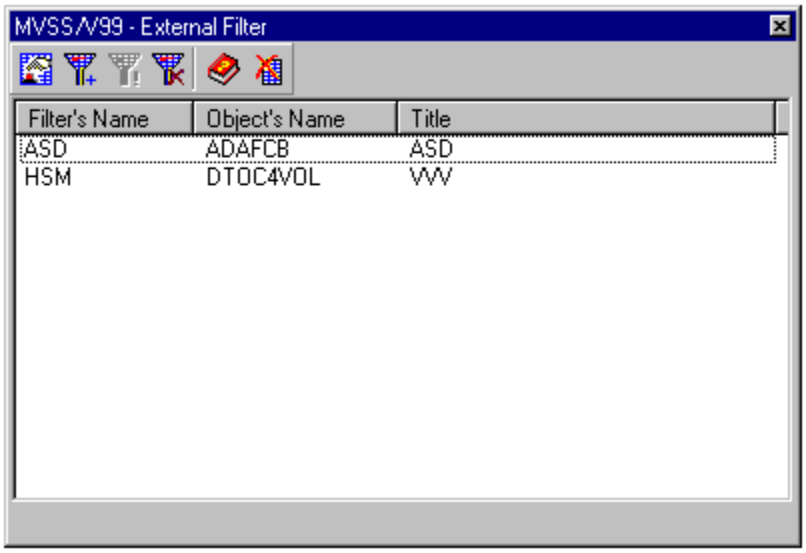

The External Filter dialog lists all the filters defined in the system.

To add a new external filter:

1. Click the **New Filter** icon.

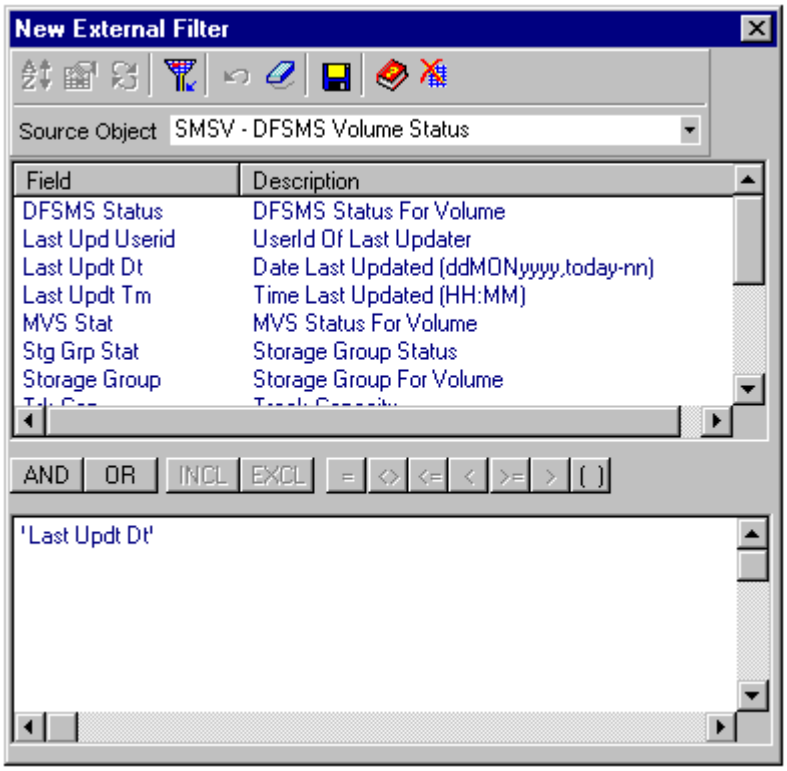

2. When the Define Filter dialog appears, select the object for which you define the filter from the **Source Object** drop-down list, then use the standard procedure to create the filter (see Filtering).
3. Click the **Save** icon on the toolbar.
4. When the Save External Filter As dialog appears, enter a name and description for the filter and click **OK**.

To edit an existing filter:

1. Select a filter from the External Filters list and click the **Edit Filter** icon.
2. When the External Filter dialog appears, use the standard procedure to edit the filter.
3. Click the **Save** icon.

Syntax of Filter Expressions

You can use wildcard characters and single quotes in filter expressions.

Use the **INCL** and **EXCL** operators to compare a character field with a test value that contains one or more pattern-matching characters. For example, the statement

```
Datasets INCL DEPT751./
```

includes all data sets whose names match the **DEPT751./** pattern. By contrast, the statement

```
Datasets = DEPT751./
```

instructs BrightStor Resource Manager to find data sets whose first-level index is exactly **DEPT751./**. In this case, the filter does not attempt any pattern matching, and is looking for a / in position 9 of the data set name. (Since IBM does not allow this on an MVS system, the statement returns no matches.)

It is a common mistake to use **DSNAMES = DEPT751./** The = operator means "exactly equal" or, in this case, data sets starting with **DEPT751** and ending with **./** in bytes 8 and 9. This exactly equal comparison is seldom what was intended. To pattern match on the **/** sign, use the pattern include operator, **INCL**.

The following examples illustrate the correct use of patterns.

To match all volumes starting with TSO, use the statement:

```
Volumes INCL TSO/
```

To match all data sets that have fixed-length records (whose record format is F, FA, FM, FBA, or FBM), use the statement:

```
Recfm INCL F/
```

When no pattern characters are specified in the test value, you can use either **INCL** or the **=** operator with the same result. The following statements are equivalent:

```
Dataset INCL an.exact.name
Dataset = an.exact.name
```

Using the **=** operator in the above example is more efficient and produces a faster search.

INCL and EXCL Lists

To provide more than one pattern at a time, enclose the list of patterns within parentheses. For example:

```
Dataset INCL (AB/, DEPT751./, SYS4/)
```

Numeric Test Values

You can enter exact numeric test values or append values with one of the following suffixes:

K, KB, KBB,

M, MB, MBB,

G, GB, GBB,

T, TB, TBB.

To allow quick and easy calculation of sizes using decimal multiplication, BrightStor Resource Manager uses decimal-based definitions for K and KB, M and MB, G and GB, and T and TB within filter expressions:

nn and nnB = nn

nnK and nnKB = nn(1,000)

nnM and nnMB = nn(1,000)(1,000)

nnG and nnGB = nn(1,000)(1,000)(1,000)

nnT and nnTB = nn(1,000)(1,000)(1,000)(1,000)

For example, you can enter the following expression:

```
'Allocated Space'  >  5GB
```

to mean 'Allocated Space' > 5000000000

However, you can use the second B in the suffix if you want BrightStor Resource Manager to use the binary based definitions of K, M, G and T:

nnKBB = nn(1,024)

nnMBB = nn(1,024)(1,024)

nnGBB = nn(1,024)(1,024)(1,024)

nnTBB = nn(1,024)(1,024)(1,024)(1,024)

For example, you can enter the following expression:

```
'Allocated Space'  >  5GBB
```

to mean 'Allocated Space' > 5368709120

The following example illustrates the way in which you might use the filter and sort functions to quickly recover some space if a DFSMS storage group has become full, causing jobs to abend.

✓ Specify a filter to select the data sets that belong to the desired storage group, **AND** have a Dsorg of ps, **AND** have some idle space. The filter statements might look like this:

```
SG-Name = groupxx AND
Dsorg INCL ps AND
Tracks_Idle > 0
```

In Dsorg, you can also use **po** instead of **ps**.

Filtering on Date Fields

Most date fields contain a real absolute date. If not, they should contain the BrightStor Resource Manager standard null date or its standard invalid date. You can filter for these in various ways.

✓ To find null dates, compare a date field to the word None. For example,

```
'Expire Date' = None
```

can be used to find all data sets without expiration dates. Testing against Gregorian date 00Jan1900 or Julian date 1900000 will achieve the same results.

✓ To find invalid dates, compare against the word Invalid. For example,

```
'Last Use Date' = Invalid
```

can be used. Testing against Gregorian date 99Dec2155 or Julian date 2155999 achieves the same results.

✓ To compare against real absolute dates, you can always enter a test date in either the Gregorian (ddMONyyyy) or Julian (yyyyddd) format. For example,

```
'Create Dt'  =  15Jul1999
```

More useful date comparisons, however, are often made to a date relative to the current date.

✓ To do this, use a value of

```
Today+nnn or Today-nnn
```

where **nnn** represents a number of days. These filters are especially useful in user views and automation scripts. For example, the filter expression:

```
'Create Dt'   >  Today-7
```

returns all entries created within the last 7 days.

Wildcard Characters

Any field whose value consists of character data (data set names, volumes, RECFMs, and so on) can be tested against patterns using the INCL and EXCL operators. BrightStor Resource Manager recognizes the following wildcard characters:

?	If this character is in the pattern, it matches any non-blank character within a character string. For example, SYS?.LINKLIB matches both SYS1.LINKLIB and SYS2.LINKLIB. However, when ? is used for filtering on text strings, it means *any character, including blanks* within the text.
*	If this character is present in a dsname string, a single level node is not checked (for example, A.*.B.*.SOURCE). The asterisk can be placed after significant characters in a string node to indicate that any following characters in the node are acceptable (for example, A.B*.SOURCE). Although the above is the most common usage, technically the * character in the pattern means that any character in the string is valid until the next period or end of string is reached. Therefore, it may also be used with text strings to mean that anything is acceptable up to the next period within the text, or if no more periods are present, to the end of the text. Note that the definition of the * character in BrightStor Resource Manager differs from the IBM usage. In most cases where you would use a * character in IBM products you should use the **/** character (described below) in BrightStor Resource Manager.
/	When this character is in the pattern, comparison to the input string terminates at the previous character. These are called prefix entries. If the prefix matches the input string up to the slash, the comparison is satisfied. For example, SYS/ matches all data sets whose names begin with SYS, regardless of what follows. This usage is similar to the way in which IBM products use the * character. For example, in IBM products, the pattern **DEPT751.*** finds all data sets with a first-level index of DEPT751. To obtain the same result in BrightStor Resource Manager, use the pattern **DEPT751./**.
!	When this character (English exclamation mark, X"5A") is encountered, the input is searched for a match on the characters that follow it. The pattern characters can occur anywhere in the input string. For example, the pattern !SYS1 matches any data set name that contains the SYS1 string anywhere in the name. The pattern !SYS1! matches any data set that ends in the string SYS1, regardless of what precedes that string.

Note: The **!** character is often translated to another on various international keyboards. An easy way to find the correct national characters on the View 3270 is to enter the filter window for any object, press F1 for Help, then browse the section Wild Characters for Pattern Matching.

Using Single Quotes

Most displayed fields contain either *character strings* (uppercase strings terminated by the first blank) or *numeric values* that do not allow embedded blanks. A few fields, however, contain *text strings,* such as message text, which can consist of several words of both upper- and lowercase letters.

If you need to filter text strings that contain phrases or multiple words separated by blanks, enter the test value within single quotes, for example, **'CI Summary'** or **'RC= 4'**. BrightStor Resource Manager translates the text string within the single quotes by replacing the first quote and all embedded blanks with **!** and the last quote with **/**, so that **'CI Summary'** becomes **!CI!Summary/**. You can enter this string directly, although it is easier to place a text string within single quotes.

Note that BrightStor Resource Manager generates a separate search for every word between single quotes. Blanks within the quoted test value are insignificant and are used by BrightStor Resource Manager to identify the words and the order of searches. If blanks are a significant part of the test value, they should be contained within double quotes (use the double quote character, not two single quotes), rather than single quotes, as described in Using Double Quotes.

You can always place your test values within single quotes.

- For hexadecimal, numeric, and date fields, BrightStor Resource Manager strips the quotes.
- For text fields, BrightStor Resource Manager replaces the first quote and all embedded blanks with **!** and the last quote with **/**.
- For character fields, BrightStor Resource Manager strips the quotes and replaces all embedded blanks with **!**.

Character fields in quotes are treated differently from text fields. Since short character fields (VOLSER, for example) often already consist of five or six characters (VOLSER INCL 'VOL123'), the quotes cannot be changed to pattern characters because this would make the pattern longer than the 6-byte maximum length of the field. Therefore, for character strings the quotes are stripped away. For text fields, the quotes can be changed to pattern characters because the length of the field allows it.

Note: If your test value is a string that contains the **<**, **=**, **>**, **(**, or **)** sign, place the entire test value within single quotes. When these symbols are not within quotes, they represent special delimiters within the filter string that designate comparison operators and parameter groupings. Treating them as operator or grouping delimiters is incorrect if you intend them to be part of a test value.

Using Double Quotes

If blanks are a significant part of the substring for which you are searching within a text field, specify your test value within double quotes rather than single quotes. BrightStor Resource Manager strips the double quotes and searches for an exact match on the value between them, including any blanks. For example, "CI Summary" is interpreted as "search for the 10-byte string CI Summary that includes the blank between CI and Summary". Similarly, "RC= 4" results in a search for an exact match on a 5-byte string that includes the blank between the = and the 4.

Note: If you specify test values using single quotes, and find that you are obtaining matches on more entries than desired, you probably need to use double quotes instead. Make sure to include the exact number of embedded blanks within your test values.

Sorting

Use the **Sort** function to sort the instances of the object listed in a table according to one or more object attributes.

Every object attribute in a table can serve as a sort key. In the case of data sets, for example, possible sort keys are **Data Set Name**, **VolTyp**, **Idle Bytes**, **Dso**, and others.

To invoke the Sort dialog:

✓ Select **Sort** from the object table's **View** menu or click the **Sort** button.

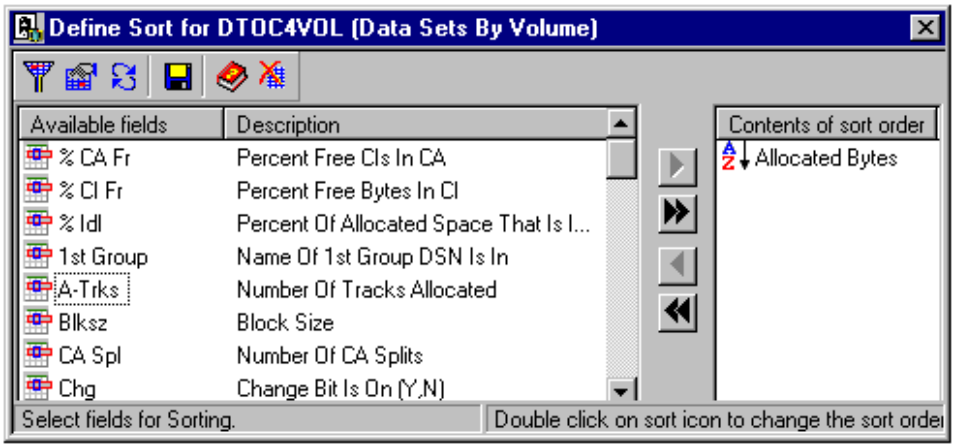

Object attributes are listed in the sort window in alphabetical order by field name. You can click at the top of the **Description** column to sort entries by description name. You can reverse the sort order by clicking again at the top of either column.

To select a sort key:

✓ Click the attribute you want to serve as a sort key, then click on the right arrow.

The selected attribute moves from the **Available Fields** to the **Contents** window. The first sort key selected (the one at the top of the list in the **Contents** window) is the primary sort key.

Note: You can define up to four sort criteria. Sort criteria defined for a given window are saved as part of a user-defined view and remain in effect until you clear the sort
criteria.

To define the sort order:

✓ Double-click the **Up** or **Down (AZ)** arrow that appears next to each attribute in the **Contents** window to toggle the ascending/descending sort order.

To remove a sort key from the **Contents** window:

✓ Select the sort key you want to remove, then click on the left arrow.

To remove all sort keys from the **Contents** window and clear the sort criteria:

✓ Click the double left arrow.

To open the Filter dialog:

✓ Click the **Filter** icon.

To execute the sorting action:

1. Click the **Set and Exit** icon.

   The dialog closes and the table from which you invoked the sort becomes active.

2. Click **Execute** to apply the sort to the table.

   - Alternatively-

✓ Click **Execute** in the Sort dialog. This function is the equivalent of clicking **Set and Exit** in the Sort dialog and **Execute** in the table.

To initiate a new data collection and then execute the sort:

✓ Click **Refresh**.

Note: If you also want to filter your data, do not click **Execute** at this point. First define your filter criteria by invoking the Filter dialog, then click **Execute** in the Filter dialog.

Sorting on Date Fields

Since BrightStor Resource Manager's standard value for a null date is Julian 1900.000, these values sort to the bottom (low value). Similarly, since BrightStor Resource Manager's standard value for an invalid date is Julian 2155.999, these most often sort to the top (high value). The only time when invalid dates do not sort to the top, occurs when fields with real dates ranging from year 2156 to 3999 are present. Dates in this high range are unlikely, but BrightStor Resource Manager supports dates through Julian 3999.365 (31Dec3999).

Note: Some objects, usually tape related, contain fields that are "date-like". Sometimes they contain real dates, but often they contain a numeric code which is assigned special meaning by the tape related software vendor. Vantage usually displays these "date-like" fields in what looks like the Julian yyyyddd (or ccyyddd) format as used by the vendor. This allows the special numeric codes to be displayed in their exact format as well, which is just the seven digit numeric value nnnnnnn. Many of these special codes are higher in value than the standard invalid date of 2155999, such as the CA-1 code for "catalog control" which is 9990000. Remember, when examining these "date-like" fields, BrightStor Resource Manager displays them as **nnnnnnn** numbers, some of which are Julian dates, others just codes.

Execute

When you select an **z/OS** object from the BrightStor Resource Manager **Object Tree** menu, BrightStor Resource Manager displays a window with a toolbar and a table that by default is empty. (In the **Data Request** panel of the object's **View Definition** function, you can request to fill the table with data when it is opened: this is the default mode for BrightStor Resource Manager objects.) To fill an empty table with data, you must click the **Execute** button. The function does not display a dialog and has no parameters.

If you just opened the table, the **Execute** command causes BrightStor Resource Manager to request the host to collect, filter, and sort the necessary data, then send the results to the PC for viewing.

If you change sort keys and click **Execute** once more, BrightStor Resource Manager redisplays the table with the new sort using data stored in memory on the host, without requesting a new collection.

If you create a new filter and click **Execute** once more, BrightStor Resource Manager initiates a new collection.

If you add a new term to the filter expression using the **AND** operator, BrightStor Resource Manager redisplays the table with the new filter using data stored in memory on the host, without requesting a new collection.

For example, if the original filter expression is:

```
Volumes INCL TSO/
```

and the revised filter expression is:

```
Volumes INCL TSO/ AND % Alloc>90
```

clicking **Execute** redisplays the table without initiating a new collection because the term following the **AND** operator is a refinement on the original expression, which has not changed. However, if the revised filter expression is:

```
Volumes INCL T/ AND % Alloc>90
```

-- or --

```
Volumes INCL TSO/ OR Volumes INCL F/
```

clicking **Execute** initiates a new collection because, in the first case, the original term of the filter has changed, and in the second, the term following the **OR** operator calls for an expansion of the filter rather than a refinement.

To force a new collection regardless of whether or not it is needed, use the **Refresh** function from the **View** menu or click the **Refresh data** button on the toolbar.

Definitions of sort keys and filters for a window remain in effect only for as long as the window is open, unless you save them with the **Save** or **Save As** commands.

Totals (z/OS Only)

The **Totals** option of the **View** menu displays the total of the values in columns containing numeric data.

To view the totals for a given object:

✓ Select **Totals** from the **View** menu.

In the totals table that appears, the rows list the numeric columns with their total values. You can print the table by clicking the **Print** button.

Statistics (z/OS Only)

The **Statistics** option displays the total, average, minimum, and maximum values for every numeric column in the table in a separate window.

To view the statistics for a given object:

✓ Select **Statistics** from the **View** menu.

In the Statistics table that appears, each row lists the numeric columns and their total, average, minimum, and maximum values. You can print the table by clicking the **Print** button.

Refresh

Use this function to force data collection on the host when it is otherwise not needed to fill a request for data. For example, if you change the sort and filter settings of a table that is already open, and if the request can be satisfied using data already in memory, a data collection cycle is not initiated after you click **Execute**. In this case, to initiate data collection on the host you need to use the **Refresh** function.

Rebuild Table

Use this function to update table properties (color and scale base) while the object is being displayed.

If you change the color or scale base of a displayed object, and want to redisplay the table with the new settings without closing and reopening the table:

✓ Select **Rebuild Table** from the **View** menu.

Monitor (z/OS only)

BrightStor Resource Manager can display the object tables in two modes: static (default) and dynamic (or Monitor). In Monitor mode the system automatically updates a specified range of records at regular time intervals.

While in Monitor mode, a snapshot of the monitored rows can be exported to a workstation log file. The setting remains in effect as long as the table is displayed in Monitor mode.

Note: When a table is displayed in Monitor mode, its **Execute** button is disabled.

To begin Monitor mode display:

1. Select **Monitor** from the table's **View** menu.

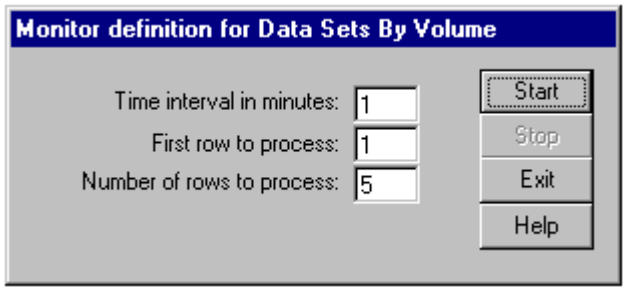

2. When the Monitor dialog appears, in the **Time interval** text box specify the number of minutes after which the display is to be updated.
3. Define the range of rows to monitor by specifying the first row and the number of rows in the corresponding text boxes.
4. Click the **Start** button. The Monitor dialog closes, and the table display is dynamically updated according to the parameters you specified.

To stop monitoring:

1. Select **Monitor** from the table's **View** menu.
2. When the Monitor dialog appears, click the **Stop** button. The Monitor dialog is closed and the **Execute** button in the table's button bar is enabled.

Settings for Monitor mode are saved as part of the user-defined views. If you save the view of the object table with Monitor mode enabled, next time you select this view the table is displayed in Monitor mode.

Monitor mode provides the basis for graphic monitoring of an object. If you start Monitor mode in the table and then define a realtime graph, the graph is updated automatically according to the time interval you specified as long as the table remains open.

Active Objects List

All active objects (open tables and dialogs) are listed in the BrightStor Resource Manager Active Objects window.

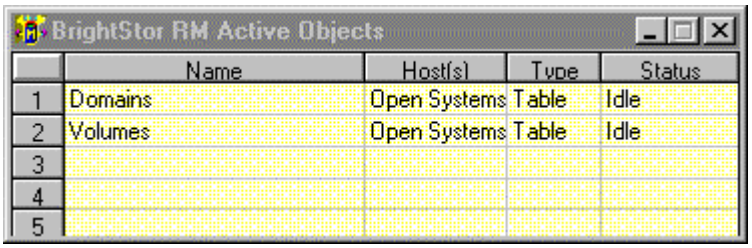

To bring the desired object to the top:

- ✓ Select the Active Objects window and click anywhere in the line or right-click and select **Bring to front**.

To close an object:

- ✓ Right-click anywhere in the line and select **Close view**.

An icon for the Active Objects list appears in the Windows taskbar.

User Views

User views reside under the source objects to which they belong. A source object contains all the fields collected by the host for that object. All fields or any subset of them can participate in user views based on that object. Source objects appear in the object tree in green. User views appear in blue.

You cannot modify or delete a source object. But you can create your own user views based on a source object. A user view consists of a subset of object attributes. Any number of user views can be derived from each source object. You can delete and modify user views as needed.

The BrightStor Resource Manager installation is delivered with a selection of pre-defined user views. You can modify and save these views under different names to create your own views.

A user view contains the following definitions:

- Data collection mode and consolidation criteria
- Filter criteria
- Sort criteria
- List of table columns
- Print settings
- HTML catalog information (if it is an HTML view)
- Graph definitions (if the view is a graph)
- Definition of collector properties (if the view is a collector)
- Definition of trend report properties (if the view is a trend report)

To define a user view:

1. Open the table (source object) for which you want to create a user-defined view.
2. Define a filter (see Filtering).
3. Define sort criteria (see Sorting).
4. Select the columns you want to display in the table.
5. Specify the type of the user view (table window, graph, and so on) and the method used to request data from the host and to consolidate it (see View Definition).
6. Define the print settings (font, headers, footers, and so on) to be used with the view (see Print Preview Tab).
7. Define a graph, if applicable to the object (see Graph Definition).
8. From the view's title menu, select **Save As**.

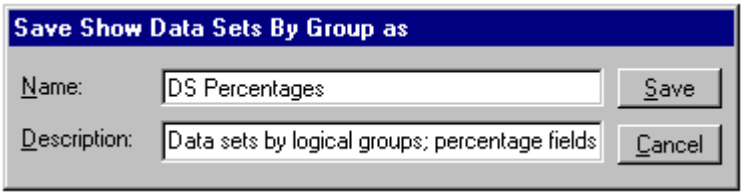

9. Enter a name and description for the new view and click **Save**.

The view you have defined is added to the list of user-defined views in the Object Tree, under the object to which the original view belongs.

Save

Use this command to save the current display view of the table, with its sort, filter, and other definitions, as the default view for the selected table. The next time you open this table, the current view appears.

The following characteristics are saved: table size and position, number of rows, number of columns, selected columns, sort and filter definitions, graph definitions, data collection mode definitions, Autowrap settings.

Save As

Use this command to save the current view as a user-defined view. When the Save As dialog appears, you are asked to assign a name to the view.

The following characteristics are saved: table size and position, number of rows, number of columns, selected columns, sort and filter definitions, graph definitions, data collection mode definitions, Autowrap settings.

You can export and import user-defined views to make them available to other users. To export a user-defined view:

1. Select **Export** from the **Tools** menu.
2. When the Export dialog appears, select the required user view from the object tree.
3. Click the **down arrow** to add the view to the **List of User Views for Export**.
4. Specify the file name under which to save the view and click **Export**.

To import a user-defined view:

1. Select **Import** from the **Tools** menu.
2. When the Import dialog appears, navigate to the desired file in **External file location** and select the view you want to import, then click **Import**.

Save Data for Demo

Use this function to save the records in a table for later viewing in demo mode. When BrightStor Resource Manager is started in demo mode, it does not connect to a host, but you can open tables that have been saved for demo.

Working Sets

A working set is a named collection of objects and views. Create working sets for quick and convenient access to your most frequently used views.

Using Working Sets

When you select **Working Sets** on the **Vantage** menu, all defined working sets are listed in the secondary menu. You can open a working set by clicking on it in the secondary menu.

To populate an empty working set:

✓ Select and drag as many objects and views as desired from the object tree to the working set window.

If other objects are present in the working set, to add a new object or group of objects to the working set:

1. From the right mouse button menu of the working set window, select **Position**, then **Add Before**, **Add After**, or **Add Under**.
2. Select the desired objects and views you want to add to the working set, then drag them to the working set window onto the object before, after, or under which you want to place them.

Objects and views are copied from the object tree with the entire hierarchy that exists below them. All the standard menu functions available for objects and views in the object tree are also available to the working set.

In the **Working Sets** tab of the Activation dialog you can specify which working sets you want to be active when BrightStor Resource Manager is started.

Defining and Customizing Working Sets

To create working sets for quick and convenient access to your most frequently used views:

1. Select **Working Sets,** then **Customize** from the **Vantage** menu or click the **Add/Delete Working Sets** icon on the main toolbar.

   The BrightStor Resource Manager Working Sets dialog appears, listing all working sets that have been defined in the system. From the right mouse menu of this dialog you can choose to view, rename, or delete the selected working set.

To create a new working set:

2. Choose **New working set** from the right mouse menu.

   A new item, named New workset, appears in the list.

3. Select the new working set and rename it.

4. Select the renamed working set and choose **View** from the right mouse menu.

   An empty window appears for the new working set.

5. To populate the working set, select and drag as many objects and views as desired from the object tree to the working set window.

Objects and views are copied from the object tree with the entire hierarchy that exists below them. All the standard menu functions available for objects and views in the object tree are also available to the working set.

Edit Menu

Use the **Select All Lines** and **Deselect All Lines** options in this menu to select and deselect all lines in the displayed table.

To select one row of the table or to change from one selected row to another:

✓ Click the desired row number.

To select multiple rows, but not all rows:

✓ Click on the first row, then press and hold the **Shift** key to select other consecutive rows or the **Ctrl** key to select non-consecutive rows by clicking on their row numbers.

To clear all selected rows:

✓ Select the **Deselect All Lines** option from the **Edit** menu.

To drag and drop one or more rows into another table:

1. Select the row or rows you want to drag. You can use any of the selection methods available:

   a. To select a *single* row, click on the row or enter the row number in the **Show line** box and press **Enter** to go to an invisible row.

   b. To select multiple *non-contiguous* rows, Ctrl-click on the line numbers of the lines you want to select

c. To select *contiguous* rows, click the left mouse button and drag the pointer over the numbers of the lines you want to select.

d. To select *all* the rows in a table, choose **Select All Lines** from the **Edit** menu.

2. Without releasing the mouse button, drag the mouse to the right until you see the drag-and-drop icon inside the table. (If you released the mouse and you are in the process of dragging a selection of non-continuous rows, press the **Ctrl** button before clicking the mouse once again - otherwise you lose the selection.)

3. Drag the icon to the desired location.

Actions Menu

In some views you can take actions on selected table rows. While the actions available for the various objects and views differ, the procedure for performing an action is similar for all objects.

1. Select one or more rows on which you want to perform the action.

2. From the table's **Actions** menu, select the desired action.

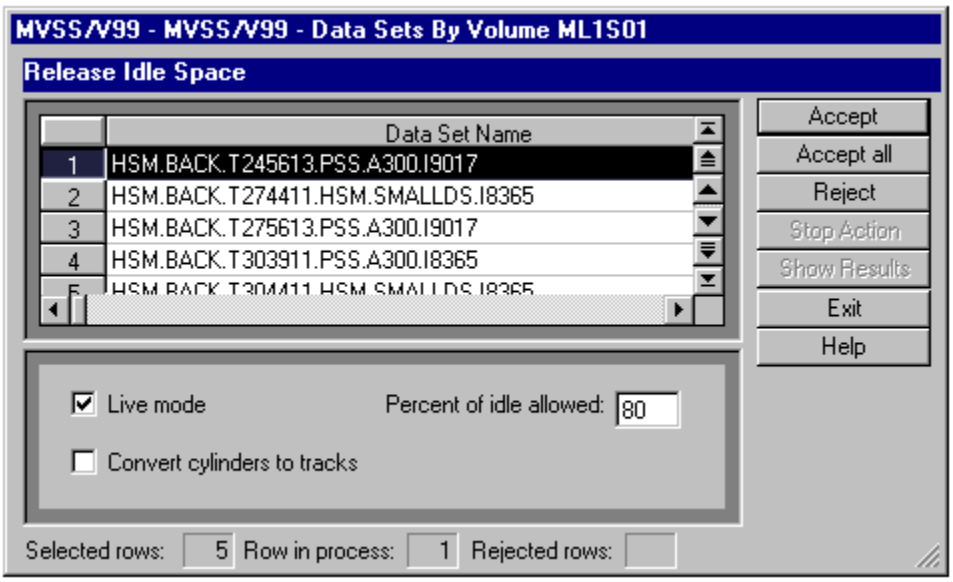

The rows you selected in the table are listed in the Actions window.

Some actions require that you supply parameters. Others ask only that you confirm the action, individually for each row or globally for all selected rows.

If more than one row appears in the Actions window, you can confirm execution of the actions row by row or globally, for all selected rows. To confirm execution of a single row:

a. Select the row and enter the desired parameter or parameters. Click **Accept**.

b. After the action is performed, the selection automatically moves to the next line. Repeat the action with the next row.

To confirm execution of all the rows from the selected one to the last row in the Actions dialog:

✓ Click Accept All.

To interrupt execution of actions confirmed with the **Accept All** command:

✓ Click **Stop**.

To reject the selected action:

✓ Click **Reject**.

To reject all selected rows and exit the dialog:

✓ Click **Exit**.

The Results pane at the bottom of the Actions window shows the results of the actions confirmed and executed by the system.

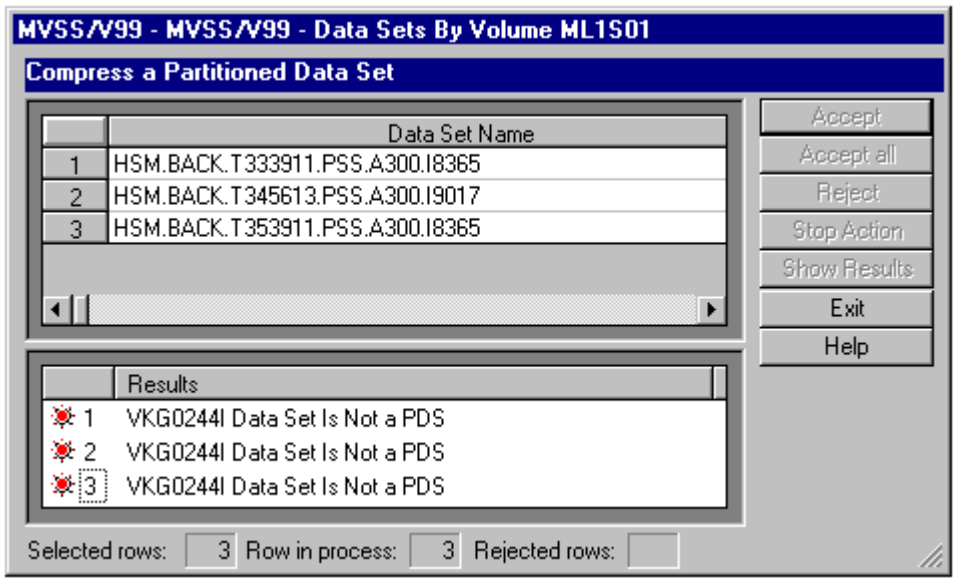

Refer to the online help for information on the specific actions available for each view.

Open

Use the **Open** function and the **Open Zoom** icon to obtain more information about a selected record. After you select **Open** from the title menu or click **Open Zoom**, BrightStor Resource Manager displays all the source objects to which you can zoom from the current table.

✓ Click on the zoom view you want to display.

Export

Use this command to output the contents of the selected table to a sequential data set on the host, to an ASCII file on the local machine, and/or to a file in a share directory. Exporting the contents of an object table to an ASCII file allows you to further process the table data using an application of your choice, such as a spreadsheet or graphic package.

Exporting data to the host is referred to as capturing. When you capture an object to a sequential data set, the entire object is written to the data set.

When exporting data to a PC file, you can select to output:

- All the rows and columns
- Only the rows and columns that are selected at the time you invoke the **Export** command
- A specified range of rows

The resulting file on the workstation is comma delimited, with all strings enclosed within quotation marks. The first line of the file contains the column headings.

To export all the rows and columns of the selected object table to a host file:

1. Select **Export** from the first menu of the table.
2. When the Export data dialog appears, click **Host File**, then click **Start**.

   The BrightStor Resource Manager log server captures the entire contents of the selected object and writes it to that object's capture file.

To export the object data to a PC file on the local machine or to a CA-Vantage Network Edition file:

1. If you want to use direct selection to specify the rows to be exported, select the desired rows before invoking the **Export** function. (If you intend to export all the rows or a range, skip this step.)
2. Select **Export** from the first menu of the table.

3. Click one of the radio buttons in the selection box to designate the rows to be exported. If you want to export the entire table, click **All rows**. If you want to export a range of rows, click the **Range** button and specify the desired range in the **From** and **To** boxes. If you selected the rows directly, click **Selected rows**.
4. Specify the location for the export file by typing in its path or by browsing to it.
5. Click **Start** to execute the export.

Print

Use this dialog to print a table:

1. Select **Print** from the title menu.
2. When the Print dialog appears, specify **Output Type** (Printer, HTML or e-mail).

   For printer output, select the desired printer and the number of copies you want printed from the appropriate drop-down lists.

   For HTML output, enter the name of the HTML file under which you want to save the table in the **File name** field.

   To attach the table to an e-mail message, click the **E-Mail** button. This selection activates Microsoft Outlook.

You can choose to print all the rows in the table, only the rows that are selected on the screen at the time you invoke the **Print** command, or a specified range of rows. To print a specified range of rows:

1. Click the **Range** button.
2. Enter starting and ending row numbers in the **From** and **To** fields.
3. Click **OK** to complete the operation.

Refer to the next section for information on how to select report data and define the print properties of your table.

Print Definition

The Print Definition dialog includes the following tabs:

- Report Data tab. Use to organize the fields in the report.
- Print Preview tab. Use to display a preview of the formatted report and change its style.

- Properties tab. Use to enter titles, apply aggregate functions to columns, and more.
- Graph Settings tab. Use to define a graph for inclusion in the report.

Saving a Print Definition

You can save the print settings of the open table as the default for this table. Next time you open the print definition dialog, it is displayed according to the default settings.

The print definition is saved as part of the user view. To save the print definition you must save the table as a user view:

✓ Close the Define Print dialog, then select **Save** or **Save As** from the object's title menu.

The following message appears when you attempt to close the print definition dialog:

```
This change will be activated next time the user view is opened. Do not
forget to save your user view to keep this change.
```

If you close the table without saving the user view, your print definitions are not saved.

Printing a Table

Refer to Print Definition for instructions on printing the table.

Report Data Tab

The Report Data tab displays the data to be included in the report. You can adjust the presentation of the data.

To adjust column width manually:

✓ Drag one of the dividers on either side of the column heading.

To adjust column width automatically to match the widest entry in the column (including the heading):

✓ Double-click the column heading.

To move a column:

✓ Drag the column heading to a new location in the table.

Grouping

You can divide a report into groups based on a common value of an attribute. For example, you can group entries in a volume space usage report by their device type, so that all 3380, 3390, and other devices are listed in separate groups. Summaries of all numeric data are calculated for each group (in addition to grand totals, which are calculated for the entire report).

You can further subdivide each group into subgroups, using common values of yet another attribute. For example, you can divide each group of devices by model (based on the **Model** attribute). Then, you can further subdivide each model by number of cylinders, and so on.

To group records by a given attribute:

✓ Drag the column heading of the desired attribute to the area above the table.

The column moves to the leftmost position in the table. The value of the grouping attribute for all entries within a group is the same. Data can be sorted and totals calculated for each group.

To create subgroups:

1. Drag an additional column heading to the area above the table.

   The column moves left to reflect the change.

2. Repeat for each subgroup.

To move a subgroup:

✓ Drag it to a new location in the area above the table.

The column moves to reflect the change.

To sort data in a group:

✓ Click the column heading of the attribute by which you want to sort. Click again to sort in the opposite order. When data is sorted by more than one attribute, the leftmost attribute (column) serves as a primary key; the precedence of secondary keys is from left to right.

Print Preview Tab

The Print Preview tab shows the results of your selections and formatting. The following controls are available:

Control	Description
Preview Page	Enter the number of the page that you want to preview or use the arrows to move the page into view.
Zoom	Drag the slider to increase or decrease the magnification.
Style	Select a preset format, which defines the font size and type, color, and other settings for the report.

Properties Tab

The Properties tab contains groups of controls used to customize the formatting and the properties of the fields in the report:

- ✓ Use **Titles** to specify the report title, page header, and page footer.
- ✓ Use **Columns** to specify which attributes (columns) to include in the report, and to specify the heading, the sort order, and the aggregate function applied to the column data.
- ✓ Click the appropriate button to select **Table** (show the data in tabular form) or **Detail** (show each record in a separate page) as the report type.

Titles

To specify the title that appears on the first page of the report:

- ✓ Select **Report title** from the **Type** list box, enter the title text in the **Report Title** box, and select the alignment by clicking the appropriate radio button in the **Alignment** box.

Page header and footer text consists of left, center, and right parts.

To specify the header that appears on every page of the report:

- ✓ Select **Page Header** from the **Type** list box, and enter the header text you want to appear in the **Left**, **Center**, and **Right** boxes.

To specify the footer that appears on every page of the report:

- ✓ Select **Page Footer** from the **Type** list box, and enter the footer text you want to appear in the **Left**, **Center**, and **Right** boxes.

You can include the following codes in any of the titles:

Code	Description
&p	Page number
&d	Date
&t	Time

Columns

The attributes of the selected object for which the report is generated are shown in the **Columns** list. Selected attributes are included in the report.

To include all the attributes in the report:

✓ Click the **Check all** button.

To exclude all the attributes, except those that have already been grouped:

✓ Click the **Uncheck all** button.

To selectively include attributes in the report:

✓ Uncheck the box next to attributes that you do not want included in the report.

Note: You cannot clear the checkbox next to an attribute that has already been included in a group (see Grouping).

For each selected (checked) attribute, you can define the column heading, the sort order, and the aggregate function.

Control	Description
Column Heading	Text of the column heading if different from the attribute name. Must be unique in a report.
Sorting	Sort order (descending, ascending, or none)*.
Aggregate function	Aggregate function to be applied to the selected column. It can be None (default) or Count.

*Sort order can be defined also by clicking the group/subgroup in the Data tab, but the **None** option is available only in the Properties tab.

Note: Aggregate functions are not applied to attributes that are used to group the records of the report.

Graph Settings Tab

Use the Graph Settings tab to define a graph for inclusion in the report.

1. Check the **Add graph to report** box to include a graph at the end of the report.

   If this box is unchecked, no graph is appended to the report, but the definitions you make on this tab are saved for later use.

2. From the **Category (X) axis** list, select the object attribute on which you want to generate the graph.

3. In the **Number of Rows** box enter the number of records you want included in the report.

4. The **Value (Y) axis** shows the list of values associated with the selected attribute. Check the boxes next to the values you want shown in the report.

   The selected attribute appears on the X-axis of the default graph.

   Each record you select in the **Number of Rows** box produces a group of traces on the X-axis of the graph; each group contains as many traces as you have selected values from the **Value (Y) axis** list.

In the graph window, a sample graph shows the result of your selections. The graph is based on sample data from the object table.

The toolbar at the top of the Graph Settings tab makes available several graph manipulation tools.

Graph Toolbar

The toolbar at the top of the Graph Settings tab makes available the following graph manipulation tools:

- **Copy** the graph to the clipboard as a bitmap.
- Change the gallery **type** (the type of graph: line, bar, area and so on).
- Toggle between **2D** and **3D** views.
- **Rotate** the graph (3D views only).
- Display the values in **Z-clustered** series (3D views only).
- Show/hide the graph **legend**.
- Show/hide the **series legend** of the graph.
- Show/hide the **vertical grid**.
- Show/hide the **horizontal grid**.

- Change the text **fonts** of the graph. After you click the icon, select the text you want to format. When the Font dialog appears, select font and size, then click **OK**.
- Edit the top, bottom, left, and right **titles** of the graph.
- Change **graph options**. Use the multi-tab dialog to format the graph and specify general graph features (General tab), to independently customize the settings for each series of data (Series tab), to customize X- and Y-axis settings (Scale tab), to set X- and Y-axis angles for viewing the graph in different perspectives (3D View tab), to specify the titles that appear on the top, bottom, left, and right sides of the graph window (Titles tab).

The Graph Settings provides balloon help for every icon in the toolbar. Refer to General Tab for more information on how to define chart properties.

Graph Definition

BrightStor Resource Manager displays this dialog when you select **Graph Definition** from the **View** menu. The Define Graph dialog appears also in response to selecting **Graph View** from the **View** menu, in case the graph has been already defined for the table.

Use the **Graph** function to display a selected set of data (rows and columns of the table) in a graph format (a format that represents numeric data as a chart). Graphs can be static or dynamic. See Types of Graphs for a discussion of the different types of graphs you can create with BrightStor Resource Manager.

To define a graph:

1. Open a table for which you want to display the graph.
2. Define a filter, sort order, and other parameters (optional).
3. Optionally, select the rows you want to show in the graph. If you do not select any rows, the graph shows all the rows visible in the table at the time the **Graph** function is invoked.

4. Select **Graph Definition** from the **View** menu or click the **Graph** button.

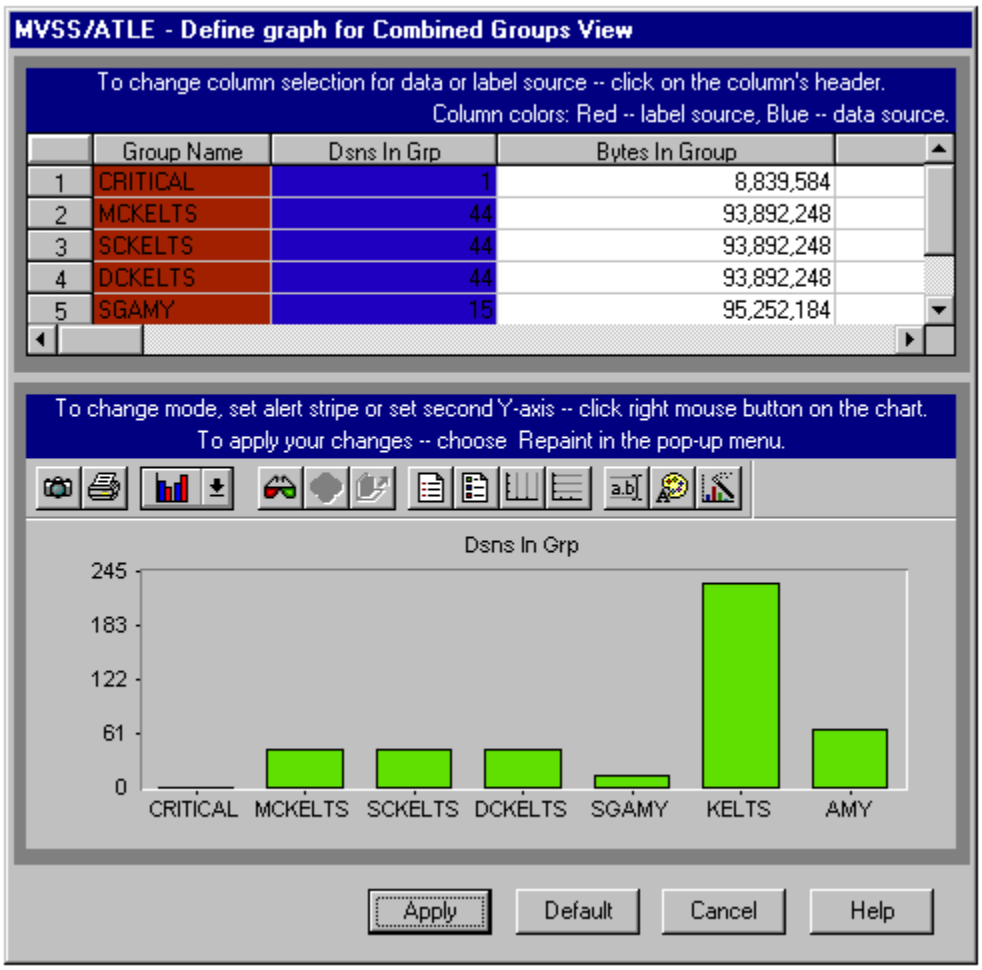

The top portion of the Define Graph dialog shows the table rows on which the graph will be based. If you selected any rows before invoking the **Graph** function, these rows now appear in the top portion of the dialog. Otherwise, this portion contains the rows visible in the table at the time the **Graph** function was invoked.

When the Define Graph dialog appears, two columns are selected: the first text column for X-axis labels and the first numeric column for data. You can change these selections as described below.

The bottom portion of the Define Graph dialog shows the graph that results from selections made in the top portion of the dialog.

Selecting X-Axis Labels

You can select only one text column for labels. Selected Text columns appear in red. To change the selected text column:

1. Right-click on the column you want to select.
2. From the pop-up menu, choose **Set column as X-axis label source**.

   The previously selected text column is automatically unselected.
3. Click **Apply** to save the graph definition and show the graph in a separate window.

Selecting Numeric Columns

You can select one or more numeric columns for the data of your graph. Selected numeric columns appear in orange. To unselect a selected numeric column:

✓ Right-click on the column you want to unselect and choose **Unselect the column** from the pop-up menu.

To select a numeric column for inclusion in the graph:

1. Right-click on the column you want to select and choose **Set column as data source** from the pop-up menu.
2. Click **Apply** to save the graph definition and show the graph in a separate window.

The chart toolbar in the middle of the Define Graph dialog provides access to a set of chart manipulation tools.

Chart Toolbar

The chart toolbar in the middle of the Define Graph dialog makes available the following chart manipulation tools:

- Copy the graph to the clipboard as a bitmap.
- Print the graph.
- Change the gallery type (the type of chart: line, bar, area and so on).
- Toggle between 2D and 3D views.
- Rotate the graph (3D views only).
- Display the values in Z-clustered series (3D views only).
- Show/hide the graph legend.
- Show/hide the series legend of the graph.

- Show/hide the vertical grid.
- Show/hide the horizontal grid.
- Edit the top, bottom, left, and right titles of the graph.
- Change the text fonts of the graph. After you click the icon, select the text you want to format. When the Font dialog appears, select font and size, then click **OK**.
- Change chart options. Use the five-tab dialog to format the graph and specify general chart features (**General** tab), to independently customize the settings for each series of data (**Series** tab), to customize X- and Y-axis settings (**Scale** tab), to set X- and Y-axis angles for viewing the chart in different perspectives (**3D View** tab), to specify the titles that appear on the top, bottom, left, and right sides of the graph window (**Titles** tab).

The dialog provides balloon help for every icon in the toolbar.

Types of Graphs

BrightStor Resource Manager graphs can be static or realtime (dynamic). Static graphs acquire data from the table once, when the graphic function is invoked. Realtime graphs acquire data periodically from a table displayed in Monitor mode (see Mode (z/OS)). Each time the table data is refreshed, it is passed to the graph window, which is also updated.

There are two modes of realtime graph display:

- Showing the latest values of one or more parameters updated periodically
- Showing the values of various parameters as they change over time. You can use this mode when the X-axis represents time and you want to update the chart with the changes that take place over time.

BrightStor Resource Manager offers approximately 30 graph types (including 2D and 3D views) to present your data in a variety of formats. Some of the available graph types are: Line, Point, Spline (Curve), Bar, Pie, Doughnut, Polar, Cube, and Conic. You can specify the graph type and its format to obtain the kind of display you prefer. You can preview the general look of a graph while defining its various features.

All the graphs can be displayed in 2D or 3D views and support special 3D effects, zoom capabilities, and rotation.

Right Mouse Button Menu

The right mouse button provides access to the following functions:

- Graph mode
- Set alert region
- Secondary Y-axis
- Repaint graph

Use the **Graph mode** function to choose between static and realtime modes. To select the graph display mode:

1. Select **Graph mode** from the right-click mouse menu, then select either **Static** or **Real**.

   Note: You can display realtime graphs only for tables in Monitor mode.

   If you chose a realtime graph, you must select the values you want shown from the secondary menu:

2. Choose **Show last samples** to show the latest values of one or more parameters over time. Choose **Show accumulated samples** to show how the current values of various parameters change over time (in this case the X-axis is the time axis).

   If you chose accumulated samples, a window appears requesting that you specify the maximum number of consecutive measurements presented at any time.

3. Enter the maximum number of measurements in the **Show samples** field.

Note: To show accumulated samples, you can have no more than one column of numeric data selected as the data source.

Use the **Set alert** region to specify marginal parameter values to provide a visual indication of parameters that fall outside of set limits.

✓ Select Set alert region from the right mouse button and enter From and To values for the alert region.

Use the **Secondary Y-axis toggle** to specify two scales for the Y-axis. This function is available only if you selected exactly two numeric columns for your data source.

Use the **Repaint graph** function to apply the current selections to the graph displayed at the bottom of the window.

View Graph

To display the graph view of the table data:

✓ Define the graph function and click **Apply** in the Define Graph dialog.

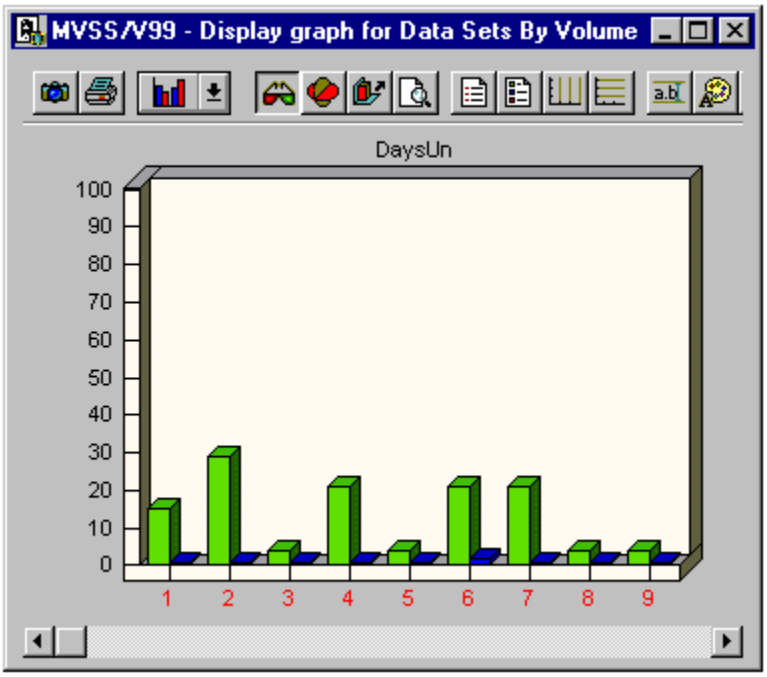

If the graph has been already defined:

✓ Select **Graph**, then **Graph View** from the table's **View** menu or click the **Graph** button.

You can modify the displayed graph view using several tools that can be accessed from a menu bar and a toolbar. The toolbar functions are identical to the functions accessed from the Define Graph dialog.

To start a realtime graph after the data has been refreshed in the object table (rows that are monitored are determined by settings in the Monitor dialog):

✓ Click the **Start** button on the graph window.

Each time the object table is refreshed, the graph is updated as well until you click **Stop** or close the graph window.

To present your data in the best graphic format, you can manipulate a large number of variables available from the toolbar of the graph window. The following tools are available to format the graph:

- Copy the graph to the clipboard as a bitmap
- Print the graph
- Change the gallery type (the type of chart: line, bar, area and so on)
- Toggle between 2D and 3D views
- Rotate the graph
- Display the values in Z-clustered series
- Zoom
- Show or Hide the legend
- Show or Hide the series legend
- Display/hide vertical and horizontal grid
- Edit the top, bottom, left, and right titles
- Change the fonts of the titles
- Change chart options

The dialog provides balloon help for every control button in the toolbar.

General Tab

Use the settings in the **General** tab to select the chart type and apply formatting features to the entire chart, including all its series.

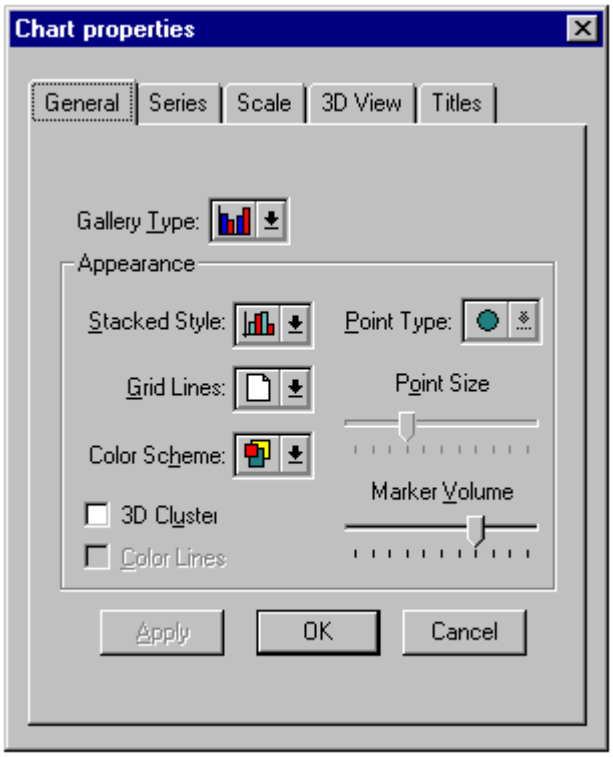

- **Gallery Type** - Click the drop-down box to select a chart type.
- **Stacked Style** - Stack the values in the chart (for Horizontal bar, Columns, and Area charts).
- **Grid Lines** - Select the grid style of the chart.
- **Color Scheme** - Select **Solid Colors**, **Colors**, and **B&W patterns**.
- **3D Cluster** - Activate/deactivate 3D-clustered charts (available only in 3D mode).
- **Color Lines** - Make 2D lines appear in the same color as the series (for Polar, Spline, and Line charts).
- **Point Type** - Select the marker shape (for Marks, Scatter, Fit to curve, Polar, and Pareto charts).
- **Point Size slider** - Select the marker size (for Marks, Scatter, Fit to curve, Polar, and Pareto charts).

- **Marker Volume** slider - Control the volume that every marker occupies; sliding to the right increases marker volume. For example, for charts of the Columns type, moving the slider to the extreme right position causes the bars to be joined (for selected charts only).

Series Tab

Use the **Series** tab to independently customize the settings for each series of data. When more than one object parameter is displayed in the same graph, you can assign different series to different parameters by checking the **Multiple Types** box and selecting the appropriate graph type for every parameter from the **Gallery Type** combo box.

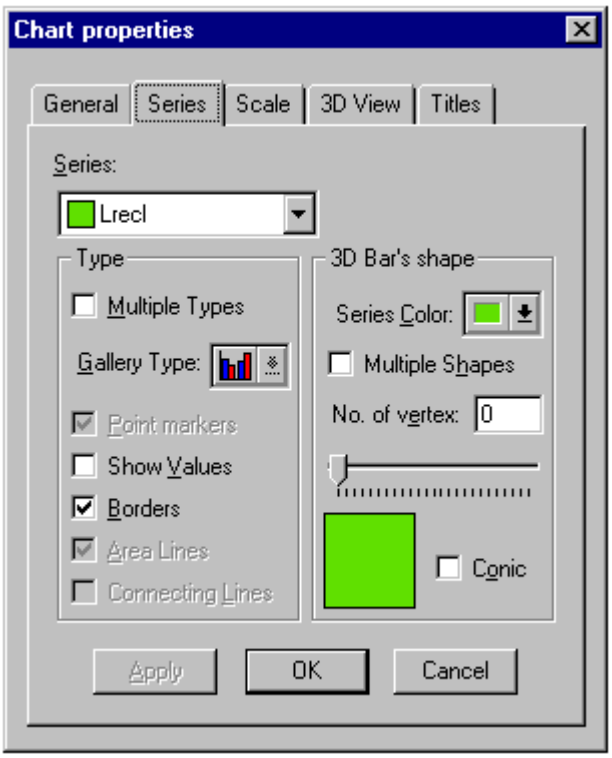

Note: When the **Multiple Types** or **Multiple Shapes** box is checked, the settings in this tab apply only to the series selected in the **Series** combo box.

- **Series** - Select the series to which you want to apply the changes.
- **Multiple Types** - If your chart contains more than one series of data, checking this box lets you select a different type of chart for each series. After you check the Multiple Types box, click the **Series combo** box to select a data series, then the Gallery Type combo box to apply one of the available chart types to the selected series.

- **Gallery Type** - Select a chart type for the selected data series (only when the Multiple Types box is checked).
- **Point Markers** - Show the points in the chart. In a Multiseries chart, to show the points for only one series, check the **Multiple Types** box without changing the chart types and check the **Point Markers** box for the appropriate series.
- **Show Values** - Show the appropriate value above the marker (for all chart types).
- **Borders** - Display the borders surrounding the markers. This feature is useful when defining conic or cylindrical charts.
- **Area Lines** - Activate/deactivate lines in area charts.
- **Connecting Lines** - Activate/deactivate connecting lines (for Polar and Scatter charts).
- **Series Color** - Select the color for the series currently displayed in the **Series combo** box.
- **Multiple Shapes** - Define different conic or cylindrical shapes in a Multiseries chart. When checked, you can select the desired series and apply a different shape to its marker.
- **No. of Vertex** - Defines the appearance of the conic and cylindric shapes (appears only for series that use these shapes).
- **3D Line Thick** - Define line thickness for 3D Line charts. You can use the slider to change the line thickness (appears only for series that use lines).
- **Base Selector** slider - Control the number of sides at the base of the cylindrical shape applied to series in Bar, Cube, and Hi-Low-Close charts.
- **Conic** - Define a conic shape for Bar, Cube, and Hi-Low-Close charts.

Scale Tab

Use the **Scale** tab to customize X- and Y-axis settings.

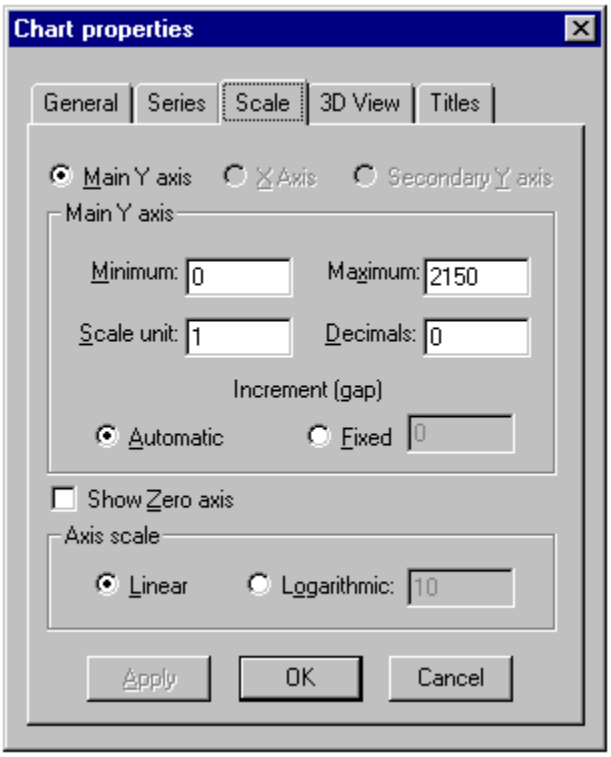

- **Axis selection** - Click the axis that you want to customize (X-axis settings are available only to Scatter charts).
- **Minimum and Maximum** - Define minimum and maximum values for the selected axis.
- **Scale Unit** - Change the automatically selected scale unit. Use this option when displaying large numbers. For example, with a scale unit of 1/1.000.000, every tick of the axis represents 1 million units.
- **Decimals** - Specify the number of digits that appear after the decimal point if the scale contains decimal numbers.
- **Automatic or Fixed gap** - Change the automatically selected granularity for the axis. For example, if the maximum value in the Y-axis is 100, and you want 10 intervals displayed, click the **Fixed** button and enter a value of 10 in the text box.
- **Show Zero Axis** - Show the zero axis when the chart displays negative values. The markers for positive and negative values are shown on different sides of the zero axis.
- **Linear or Logarithmic** - Select a linear or logarithmic scale. If you select a logarithmic scale, you can provide the base (base 10 is the default).

3D View Tab (Rotation Dialog)

Use the **3D View** tab to rotate the chart. You can set X- and Y-axis angles to view the chart in different perspectives.

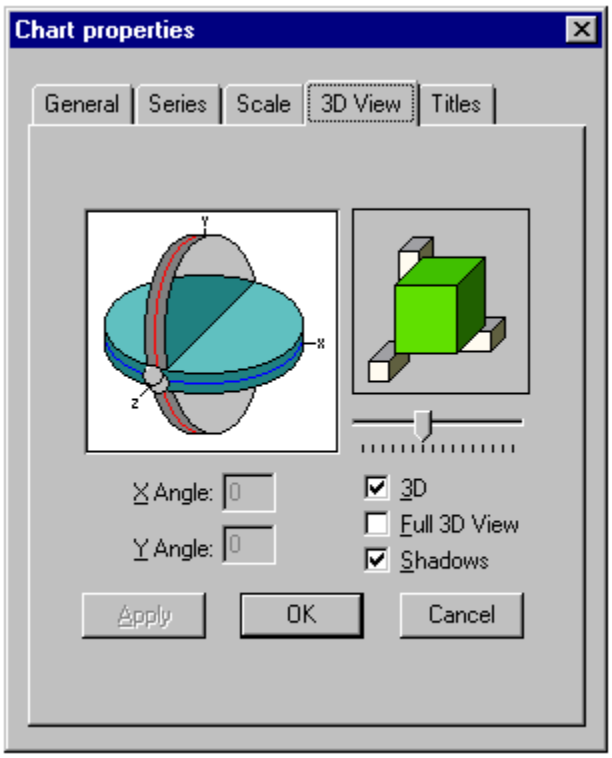

To rotate the chart:

1. Check the **Full 3D View** box.
2. Drag the marbles to the desired angles or provide the angle values in **X Angle** and **Y Angle** fields (the resulting view is shown in the sample rotation frame).
3. Click the **Apply** button to modify the angles in the displayed chart while remaining in the 3D View dialog, or click **OK** to apply the angles and exit the 3D View dialog.

To turn off rotation:

1. In the 3D View dialog, uncheck the **Full 3D View** box.
2. Click **OK** to exit the dialog.

To control the 3D depth of the chart:

✓ Use the sliding control located under the graph preview. Dragging the sliding control to the right increases the chart's depth (in comparison to its width); dragging the slider to the left decreases the chart's depth.

To control the 3D shadows:

✓ Check the **Shadows** box to turn on shadows for markers in 3D mode (default setting). Unchecking the **Shadows** box causes the back side of the markers to assume the color assigned to the front of the marker.

Titles Tab

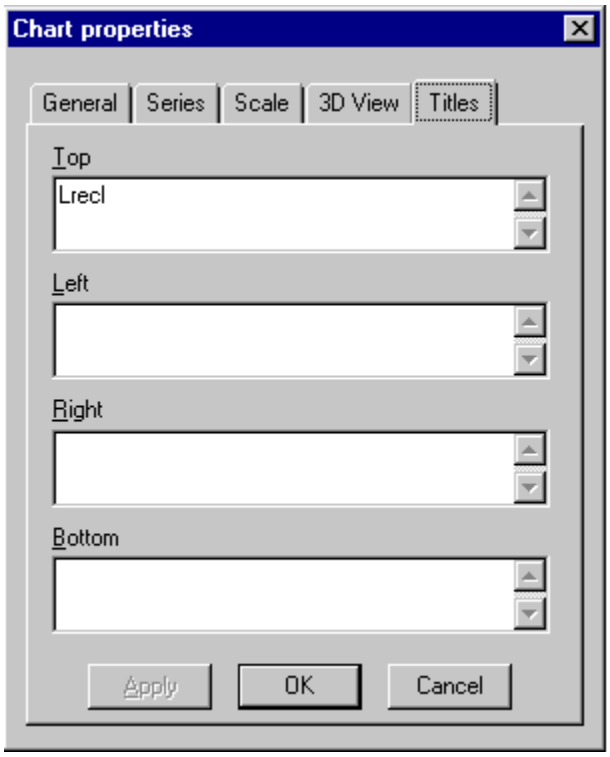

Use the **Titles** tab to specify the titles that appear on the top, bottom, left, and right sides of the graph window. Enter the desired titles in the appropriate fields and click **Apply**.

Saving Graph Definitions

Graph definitions are saved as part of the default or user-defined view of the table.

To save a graph definition:

1. Design a graph using the Define Graph dialog.
2. Display the graph by clicking **Apply**.
3. Size and position the graph on the screen as desired.

4. Change the appearance of the graph using the tools available in the toolbar to accommodate the best presentation of the data.
5. Close the graph by clicking **Close**.
6. From the object menu, select **Save** (for the default view) or **Save As** (to save it as a separate user-defined view).

You can use the saved settings to display a new graph based on different data from the table:

1. Select any number of rows in the object table.
2. From the **View** menu, select **Graph** or click the **Graph** button.

   The graph you defined earlier is redisplayed based on the data in the selected rows.

Starting a Realtime Graph

To display a realtime graph that is updated dynamically as the data changes, you must synchronize the graph with the **Monitor** function.

1. Open the desired object and adjust table height so that it shows only the number of rows you want included in the graph.
2. From the **View** menu, select **Graph**, then **Graph Definition**. Set the graph parameters as desired and click **Apply**.
3. From the **View** menu, select **Monitor**. Set the Monitor parameters as desired, and click **Start**.

   When the data is refreshed in the object table, the graph is updated as well. The rows that are monitored are determined by a setting in the Monitor dialog.

You can stop the realtime graph anytime by clicking the **Stop** button in the Monitor dialog or by closing the table.

Example 1: Monitoring Free Cylinders of Five Largest Volumes

Begin by preparing the data in the Dasd Volume Online table:

1. From the Object Tree, select **Dasd Volumes Online**, then **Space and Other Attributes**. Adjust the window to show five or six rows.
2. Click the **Sort** button and sort the table in descending order by **Device Capacity in Cylinders** (Cyls).

Generate the graph as follows:

1. Click the **Graph** button.
2. On the table displayed at the top of the definition form, select **Volume** as the label source and **FrCyls** as the data source.
3. Click the **3D toggle** on the toolbar to display a 3D graph.
4. Click the **Series Legend** button on the toolbar to have the legend show on the right side of the chart. The legend identifies the color of each column in the display.
5. Click the **Horizontal Grid** button to display the horizontal grid lines.
6. Click the **Chart Options** icon, select the Edit Titles tab, and enter

   ```
   Monitor Free Cylinders of Five Largest Volumes
   ```

   in the **Top** text box and "Alert strip: 0-500 free cylinders" in the **Bottom** text box.
7. Click the **Apply** button.

To start monitoring:

1. From the **View** menu, select **Monitor**. Set **Time interval** to 1 minute, **First Row** to 1, and **Number of Rows** to 5.
2. Click **Start**.

   When the data is refreshed in the object table, the graph is updated automatically every minute for as long as the table remains open.

You can use the resulting graph as a permanent monitor on your screen. You can minimize the table on which the graph is based, and the graph continues to be updated every 60 seconds. If, however, you close the table, the system ceases to update the graph.

Example 2: Showing % Allocated vs. % Idle for Group Space Quotas

The aim of the following graph is to show the ratio between allocated and idle space for groups that have the largest percentage of idle space. Since both values in this graph are percentages, the ratio between the two values can be shown by stacked bars that together amount to 100 percent.

To prepare the data in the table:

1. Sort the Group Space Quotas table in descending order by % Idl.
2. Adjust the size of the table to show five records (rows).

To generate the graph shown in this example:

1. Click the **Graph** button.
2. Select **Group** as the label source.

3. Select % **Allo** and % **Idl** as data sources.
4. Click **Apply**.
5. After the graph appears, click the **3D toggle** to display a 3D graph.
6. Click the **Chart Options** icon. On the **General** tab, click the **Stack Style** drop-down button and select **Stack** 100%.
7. Click the **Series Legend** icon to display the legend on the right side of the chart. The legend identifies the color of each column in the display.

Trend Reports (z/OS only)

You can view the history of selected storage management objects by graphically presenting trend data accumulated on the host in the BrightStor Resource Manager log files. The program extracts data about selected BrightStor Resource Manager objects over time and stores it in a database on the Windows Client. This operation is performed by collectors. Based on this data, the program generates reports that graphically display the changes selected objects underwent during the collection period. The program can also perform an analysis of the accumulated trend data and project it into the future for any desired period of time. This feature is especially important for capacity planning.

Both collectors and reports appear in the object tree: collectors under the object whose data they collect, reports under the collector on which they are based.

Note: Do not confuse the Trend Reports described here, which are based on data extracted by collectors and accumulated in a local database, with the Light Trends generated directly from log data accumulated on the host.

To define a new collector or modify the properties of an existing collector:

1. Select a source object on the object tree and click its plus sign (**+**) if one is present.

   If a collector has already been defined for the object, it appears directly under it.

   To modify its definition:

2. Select the collector and choose **Collector Properties** from the right mouse menu (see Collector Properties).

   If no collector has been defined for the object, you can create a new one.

   To do so:

3. Select the source object for which you want to create the collector and choose **New Collector** from the right mouse menu.

Each trend report is attached to a collector and resides on the object tree below the collector that collects data for it.

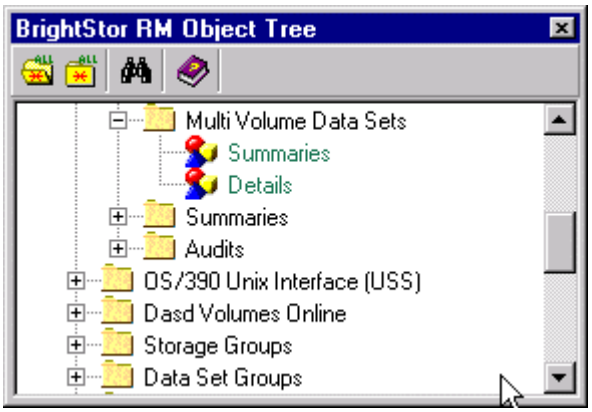

To define a new report or modify the properties of an existing report:

1. Select a collector object on the object tree and click its plus sign (**+**) if one is present.

   If a trend report has already been defined for the collector, it appears directly under it. To modify its definition:

2. Select the trend report and choose **Trend Report Properties** from the right mouse menu (see Trend Report Properties).

   If no Trend report has been defined for the collector, you can create a new one.

   To do so:

3. Select the collector under which you want to create the trend report and choose **New Trend Report** from the right mouse menu.

You can define a separate collector for each object, and associate as many reports as needed to the collector.

Data Collection

Data collection for trend reports takes place in two stages:

- Logging host object data into BrightStor Resource Manager log files. This operation is carried out on the host.
- Retrieving data from the BrightStor Resource Manager log files into the trend database. This stage is performed by the Collector on the client.

Defining Logs in BrightStor Resource Manager

You can instruct BrightStor Resource Manager to log the contents of an object to a sequential data set on the host. For each object, BrightStor Resource Manager maintains a separate log file. Data from successive logging operations is appended to the same log file.

Retrieving Log Data into the Trend Database

Collectors periodically retrieve selected data from the BrightStor Resource Manager log files on the host, and store it in the trend database.

The BrightStor Resource Manager log file contains a snapshot of the object for which the data is collected, based on a defined filter on selected records and all fields. The collectors can further filter this data and copy only selected fields and records to the trend database.

When instructing the collectors to retrieve host data into the trend database, you must bear in mind the logging regime that has been defined for the given object in BrightStor Resource Manager. For example, if an object is logged in BrightStor Resource Manager once a day, it is wasteful to make a collector connect 24 times to the host only to find that there is no new data in the log file.

To ensure that the data is continuous, a collector begins collecting data from the point where it stopped on the previous occasion.

Data Consolidation

The program maintains a cumulative trend database. If many objects are tracked at high frequency, the size of the trend database can grow rapidly. To be able to continue data accumulation without the database reaching unreasonable proportions, the program provides a maintenance function that consolidates the old data according to user-defined rules. These rules are defined in the maintenance policy.

The maintenance function replaces all the records accumulated over a specified period of time with three records, containing minimum, maximum, and average values.

When defining a maintenance policy, you specify a grace period (between the present and a designated date in the past) during which no consolidation is performed. For the period preceding the grace period, the maintenance policy specifies a unit of time for which all the records are consolidated into a set of three records (minimum, maximum, and average). For example, if you select "every three months," all the records accumulated over each three-month period are consolidated into a set of three records.

Note: If fewer than three records have been collected during the selected consolidation period, the maintenance operation is not executed.

Collector Properties

You can reach the **Collector Properties** function from the right mouse menu. The Collector properties definition dialog contains the following tabs:

- Hosts
- Fields
- Maintenance
- Schedules
- Advanced

To define a new collector or modify the properties of an existing one:

1. If this is a new collector, enter a name for it in the **Name** field.
2. Click the arrow next to the **Beginning date...** field. Use the calendar that appears to select a date on which to begin data collection.
3. Click each tab in turn and make the appropriate selections.
4. After you complete all the definitions, click **OK**.

The new collector appears on the object tree under the appropriate object.

Note: You can activate the collector manually, by clicking on it in the Object Tree, or automatically, by the system's main Scheduler.

Hosts

Use the Hosts tab to specify the hosts on which you want the collector to collect data. The tab contains a list of hosts that have been defined on the Windows Client.

- Click **All hosts** to collect data from all the hosts in the list.
- Click **Selected hosts** to collect data from hosts selected on the client.
- Click **Connected hosts** to collect data from hosts to which the client is connected.
- Click **Specific hosts,** then check the box next to the host from which you want the collector to collect data.
- Check the **Show selected hosts only** to display in the hosts list only hosts from which the collector collects data.

Fields

Use the Fields tab to specify the object fields from which the collector will collect data. The name of the source object is listed at the top of the tab, and the **Available Fields** list shows the fields of the source object.

When the tab is first displayed, the object name and the first numeric field in the table are selected by default. You can change these selections by adding and removing fields as needed.

1. Select the fields on which you want to collect data and click the right arrow to move them to the **Selected Fields** list. To move all the fields at once, click the double right arrow.

   You can remove fields from the **Selected Fields** list by selecting them and clicking the left arrow. To remove all the fields at once, click the double left arrow.

   You can rearrange the order of the fields by selecting a field and moving it up or down with the up and down arrows.

2. Click the **Define filter** button to narrow the data retrieved by the collector.

   The standard filter definition dialog appears.

3. After you define the filter, click the **Verify filter on host** button on the toolbar to check the validity of the filter before saving it. Click the **Save** button on the toolbar.

Note: If you change the selected fields being used by a collector, you affect the future operation of that collector. The collector will no longer retrieve fields you have removed from the list and will begin retrieving fields you have added. This can lead to inconsistency between data retrieved by the collector before and after you changed the
definition.

Maintenance

Use the Maintenance tab to define a new maintenance policy or modify an existing one. The tab lists all the policies that have been defined for the collector. To add a new policy (or to modify an existing one):

1. Click the **Add** button (or select an existing policy and click **View**).

   The New policy dialog that appears lets you define a time slice for the consolidation policy (in minutes, hours, days, weeks or months) and a period of time that should not be affected by the consolidation policy (in minutes, hours, days, weeks or months).

2. Specify the duration for which data collected by the collector is to be consolidated into a one record set. (Consolidation compresses data collected during a specified period into a set of three records: a minimum, a maximum, and an average record.) Enter a value in the **of** field and select a unit from the drop-down list. Units can be Minutes, Hours, Days, Weeks, or Months. For example, if you want data for every period of three months to be consolidated, enter 3 in the **of** field and select Months from the drop-down list.
3. Specify the period, counting backward from the present, during which no data consolidation is to be performed. Enter a value in the **while...** field and select a unit from the drop-down list. Units can be Minutes, Hours, Days, Weeks, or Months. For example, if you want data collected during the last month not to be affected by consolidation, enter 1 in the **while...** field and select Months from the drop-down list.

Note: If during the period you selected fewer than 3 records were collected, the consolidation is not executed.

4. Click **OK** to save the definition or **Cancel** to clear and close the New maintenance dialog.

   A summary of the new policy appears in the list of policies on the tab.

You can repeat the process to add as many policies as needed.

To delete a policy from the list:

✓ Select the policy and click **Delete**.

Schedules

Use the Schedules tab to define one or more schedules for collector and maintenance operations. The tab lists, in separate lists for collector and maintenance operations, all the schedules that have been defined. To add a new schedule (or to modify an existing one):

1. Click the **Add** button in the **Schedule collector execution** section (or select an existing schedule and click **View**).
2. When the New schedule dialog appears, specify the frequency with which the schedule is to activate the collector. Enter a value in the **Every** field and select a unit from the drop-down list. Units can be Minutes, Hours, Days, Weeks, or Months. For example: "Every 10 Minutes" or "Every 12 Hours".
3. If you chose Weeks or Months as the unit, the **On** field becomes selectable. From the drop-down box, select a day of the week (Monday, Tuesday, etc.) or a day of the month (1-31), depending on the unit. For example: "Every 1 Week On Monday" or "Every 3 Months On 1".

   Note: A value of 31 always refers to the last day of the month, even in June and in February.

4. If you chose Days, Weeks, or Months as the unit, the **At** field becomes selectable. Enter a time value in this field with the format HHMM. For example: "Every 2 Weeks On Friday At 17:00".
5. Click **OK** to save the definition or **Cancel** to clear and close the New schedule dialog.

   A summary of the new schedule appears in the **Schedule collector execution** list at the top of the tab, and a **Schedule maintenance** section appears at the bottom of the tab.

   You can repeat the process to add as many schedules as needed.
6. Click the **Add** button in the **Schedule maintenance** section at the bottom of the tab and repeat the same operations to define a schedule for maintenance operations.

To delete a schedule from the list:

✓ Select the schedule and click **Delete**.

Advanced

Use the Advanced tab to specify database and data table names.

The Database field contains the name of the file that stores the data retrieved by the collector. It is recommended to use the default name: TREND.MDB.

The Data Table field contains the name of the database table that stores the data retrieved by the collector. It is recommended to use the default, which is the same as the collector name.

Trend Report Properties

You can reach the **Trend Report Properties** function from the right mouse menu. Use this dialog to define a new report or modify the properties of an existing one.

If the collector has already collected data, the top portion of the dialog shows the data used in the trend definition; otherwise the table is empty.

When the dialog appears, two columns are selected: the first text column for labels and the first numeric column for data. You can change these selections as described below.

The bottom portion of the dialog shows the chart that results from selections made in the top portion of the dialog.

Right Mouse Menu for the Top Portion

✓ Right-click on the top portion of the definition dialog to select text and numeric columns.

The right mouse button provides access to the following functions:

- Set column as legend source
- Set column as data source
- Unselect column
- Select legend values

Use **Set column as legend source** to select a text column for labels. The selected text column appears in red. By default, the first text column is set as the label source. To change the selected text column:

1. Right-click on the column you want to select.
2. From the pop-up menu, choose **Set column as legend source**.

   The previously selected text column is automatically unselected.

Use **Set column as data source** to select one or more numeric columns for the data of your trend report. Selected numeric columns appear in orange. By default, the first numeric column is set as the data source. To unselect a selected numeric column:

✓ Right-click on the column you want to unselect and choose **Unselect column** from the pop-up menu.

To select a numeric column for inclusion in the trend report:

✓ Right-click on the column you want to select and choose **Set column as data source** from the pop-up menu.

Use **Unselect column** to unselect a legend or data source.

Use **Select legend values** to select the object instances (values) you want shown in the report.

By default, each text column you select produces as many traces on the report as the number of object instances (values) available multiplied by the number of selected numeric fields. For example, if you selected **Free Space** and **Occupied Space** as numeric fields, and **Pool Name** as the text column, and if data has been collected on seven pools, the report produces fourteen traces, showing total free and occupied space on each of the seven pools for which data was retrieved by the collector.

To show only selected values of a text column:

1. Uncheck the **All values** box to show the list of available values.
2. When the values dialog appears, select the values you want included in the report. Use **Select/Deselect All** and **Show selected values only** to help you make selections when the list contains a large number of values.

   You can manually add values that are not in the list (for example, pools you expect to be defined in the future and which you want included in the report).

   To do so:
3. Enter the name of the new value and click the **Add** button.

   By default, the program creates a trace for each selected value.

   To total all the selected values and show them as one trace:
4. Check the **Total the selected values** box. Click **OK**.

Right Mouse Menu for the Bottom Portion

✓ Right-click on the bottom portion of the definition dialog to change the appearance of the chart.

The right mouse button provides access to the following functions of the BrightStor Resource Manager Manager:

- Report name
- Report heading
- Report period
- Report sensitivity
- Report output
- Report schedules
- Y-axis parameters
- Repaint chart

Use **Report name** to specify the name under which the report will appear in the object tree.

Use **Report heading** to specify the heading printed at the bottom of the report.

Use **Report period** to specify the period covered by the report.

✓ Click the appropriate button to select a relative or absolute period.

If you select an absolute period, enter values in the **Start Date** and **End Date** fields. Click the expand button to the right of the field to access a calendar that helps you make the selection.

If you select a relative period, (a) select a starting point from the **Start from** drop-down list; (b) enter a value in the **Minus** box specifying the count-back to the beginning of the period; (c) then enter a value in the **For** box and select a unit of time to specify the duration of the period.

Examples of a relative period are: "Start from now, minus 10 days, for 10 days." "Start from the start of the month, minus 24 months, for 12 months."

Use **Report sensitivity** to specify the frequency of collection points and the method of aggregation between them:

1. From the **Time Resolution** drop-down list, select the frequency with which data points are to be collected (every minute, hourly, daily, and so on.). Aggregation is not performed if you choose raw data.
2. Select the method of aggregation between collection points (Average, Minimum, Maximum, or any combination of these).

Note: Aggregation is not allowed if you chose to total selected values.

Use **Report output** to select a method to output the report. You can send the report to the display, to a printer, or write it to disk as a graphic file (.bmp or .jpg).

If you choose the printer option, you must specify the printer to which you want the report sent. If you choose to write the report to a graphic file, the report name is used as the default file name. You can specify a different file name by clicking the expand button to select or specify a path and file name for the graphic file.

Use **Report schedules** to define one or more schedules for the report. To add a new schedule:

1. Click the **Add** button.
2. When the New schedule dialog appears, specify the frequency with which the schedule is to activate report execution. Enter a value in the **Every** field and select a unit from the drop-down list. Units can be Minutes, Hours, Days, Weeks, or Months. For example: "Every 10 Minutes" or "Every 12 Hours".
3. If you chose Weeks or Months as the unit, the **On** field becomes selectable. From the drop-down box, select a day of the week (Monday, Tuesday, etc.) or a day of the month (1-31), depending on the unit. For example: "Every 1 Week On Monday" or "Every 3 Months On 1".

   Note: A value of 31 always refers to the last day of the month, even in June and in February.

4. If you chose Days, Weeks, or Months as the unit, the **At** field becomes selectable. Enter a time value in this field with the format HHMM. For example: "Every 2 Weeks On Friday At 17:00".

Use **Y-axis parameters** to specify minimum and maximum values and their units on the Y-axis, as well as the value at which the alert line appears on the Y-axis.

Use the **Repaint chart function** to apply the current definitions to the chart displayed in the lower part of the window.

Saving and Activating the Report

✓ After you complete all the definitions, click **OK**.

The new report appears on the object tree under the appropriate collector.

Note: You can activate the collector manually, by clicking on it in the Object Tree, or automatically, by the system's main Scheduler.

Example of Collector and Trend Report Definition

The following example illustrates the process of defining a collector and trend report to monitor, in this case, the growth of pools over time. The purpose of the report is to track the rate at which pool allocated space approaches pool capacity and to forecast when pool capacity will have to be increased. This is accomplished by taking periodic snapshots of allocated space and capacity for each pool, then generating a trend report based on the accumulated data.

1. Make sure that trend logging has been defined for the Space and Other Attributes object under Storage Groups (**Object Trend Logging** function under the **z/OS** menu). The log file of this object represents the data source for the trend report.
2. Define a collector:
   a. Right-click on the Space and Other Attributes object in the object tree and select **New Collector** from the pop-up menu.
   b. Set a beginning date for collection. If a new log file is created for the object every month, you can specify the beginning of month as the start date for the collector.
   c. Specify the host or hosts on which you want to collect pool trends.

d. On the **Fields** tab select the Storage Group Name, Allocated Bytes/Cylinders, Capacity Bytes/Cylinders, and Percent Allocated fields.

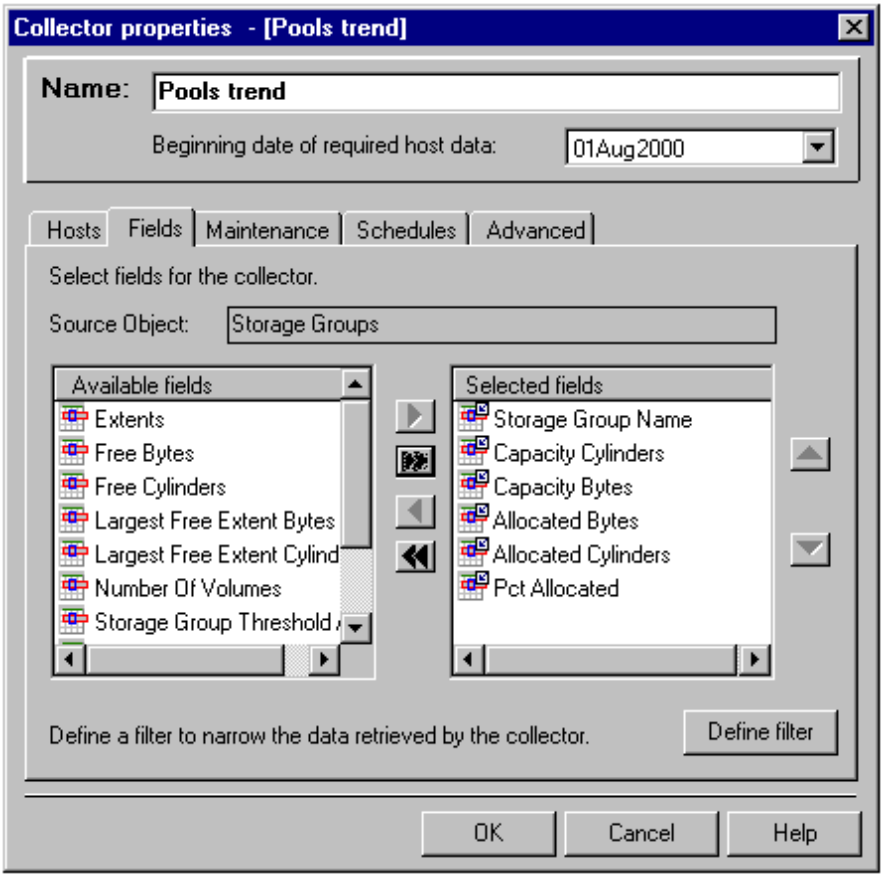

e. On the **Schedules** tab define a data collection schedule. For example, if you collect data every day at 17:00, the Scheduler automatically runs the Collector every day at 17:00 to take a snapshot of pool allocated space and capacity. The data is stored locally in a database.

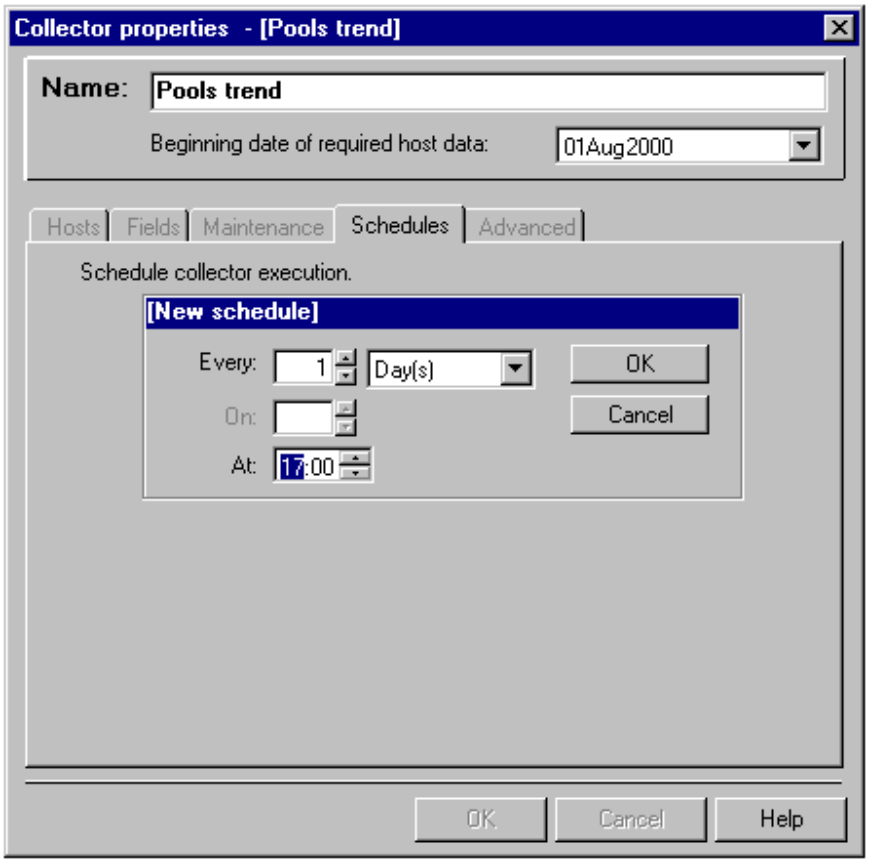

f. Click **OK** to save the schedule and again to save the collector definition.

3. To collect data from the host immediately, run the collector manually by clicking on it in the object tree.

Note: To create a trend report based on a real data sample, do not run the trend report wizard before the collector has collected host data at least once.

4. Define the trend report:

   a. In the object tree, right-click on the collector you have just defined and select **New Trend Report** from the pop-up menu.

   b. Select the Allocated Bytes and Capacity Bytes columns as the data source by right clicking on these columns and selecting **Set Column as Data Source**. Storage Name is selected automatically as the legend source. The trend report chart displays two lines for each pool: allocated space and capacity.

c. Right click on the **Host** column and choose **Select Legend Values**. Uncheck **Total the selected values** and click **OK**.

d. Right click on the chart and select Report Period. By default report covers one year from the beginning of the year. Change it as needed.

e. If you want the Scheduler to automatically run the Trend Reporter and periodically generate the report, right click on the chart, select **Report Schedules**, and add a schedule. Alternatively, run the Trend Reporter manually to generate the report.

f. Click **OK** to save the report.

5. Run the report for the first time manually to check its appearance by clicking on it in the object tree.

Light Trend Definition

You can generate trend reports directly from the object's View menu, based on log data collected by BrightStor Resource Manager on selected objects.

Note: Do not confuse the Light Trends described here, generated directly from log data accumulated on the host, with the Trend Reports based on data extracted by collectors and accumulated in the local database.

Trend reports are defined on the Log view of the object data. Perform the following preliminary steps before generating the trend report:

1. Display the object's log view by selecting **Mode**, then **Log** from the object's **View** menu. In the Log Collection Period dialog, select the range of dates for the log data.

   Note: Log data is available only if you have previously specified a logging operation for the selected object using the Log Operations wizard.

   The table is redisplayed in a new window, showing the log data of the object.

2. Filter the data to show only the records you want included in the trend report.

   When the filtered data appears, optionally, select the rows you want shown in the trend definition dialog. If you don't select any rows, all the rows present in the table at the time the **Light Trend Definition** function is invoked appear in the trend definition dialog. Note, however, that the trend report is always generated from all the records present in the object, not only those that appear in the definition dialog.

Note: To avoid skewing the data by volumes or pools that happen to be locked at the time of data collection, it is recommended to filter out volumes and pools that have a value of 0.

Use the following procedure to define the trend report:

✓ From the **View** menu select **Light Trend**, then **Light Trend Definition**.

The top portion of the trend definition dialog shows the table rows used in the trend definition.

When the dialog appears, two columns are selected: the first text column for labels and the first numeric column for data. You can change these selections as described below.

The bottom portion of the dialog shows the chart that results from selections made in the top portion the dialog.

Selecting Labels

You can select only one text column for labels. Selected text columns appear in red. By default, the first text column is set as the label source. To change the selected text column:

1. Right-click on the column you want to select.
2. From the pop-up menu choose **Set column as legend source**.

The previously selected text column is automatically unselected.

Selecting Numeric Columns

You can select one or more numeric columns for the data of your trend report. Selected numeric columns appear in orange. By default, the first numeric column is set as data source. To unselect a selected numeric column:

✓ Right-click on the column you want to unselect and choose **Unselect the column** from the pop-up menu.

To select a numeric column for inclusion in the trend report:

✓ Right-click on the column you want to select and choose **Set column as data source** from the pop-up menu.

Right Mouse Menu

✓ Right-click on the lower portion of the definition dialog to change the appearance of the chart.

The right mouse button provides access to the following functions:

- Chart parameters
- Trend titles
- Repaint graph

Use **Chart parameters** to specify minimum and maximum values on the Y axis and the value at which the alert line appears on the Y axis.

Use **Trend titles** to specify the main title and the Y-axis title that appear on the chart. Enter the desired titles in the appropriate fields and click **OK**.

Use the **Repaint chart** function to apply the current definitions to the chart displayed in the lower part of the window.

You can repeat the above procedures as many times as necessary, until you are satisfied with the appearance of the chart.

Saving and Showing the Report

✓ Click **Apply** to save the trend report definitions and show the report in a separate window.

Toolbar and Main Menu

The BrightStor Resource Manager Main Toolbar provides quick, single-click access to the main menu functions. An icon for the BrightStor Resource Manager toolbar appears in the Windows taskbar.

The main menu contains the following entries:

- **BrightStore Resource Manager** provides access to the object tree (see Object Tree Menu), host list (see Host List), consoles (see Consoles (z/OS only)), list of working sets (see Working Sets), list of active objects (see Active Objects List), Scheduler (see Scheduler), and the list agents (see Agent List).
- **z/OS** provides access to operator commands, the configuration of System and Allocation Manager Parameters (see System Parameters (z/OS only)), Member Editor (see The Member Editor (z/OS only)), the creation and management of log operations (see Logging and Capturing Data (z/OS only)), Automation Scripts, Audit Scripts (see Defining Audit Scripts), and System Scripts (see Audit and System Scripts (z/OS only)), the setup of HTML categories (see Categories), and the management of external filters (see External Filters (z/OS)).
- **BrightStor Resource Manager** provides access to functions of the BrightStor Resource Manager Monitor. Refer to the BrightStor Resource Manager Monitor Users Guide for details.

- **Tools** provides access to the BrightStor Resource Manager Activation preferences, as well as to data export and import functions.

In addition, the Toolbar contains the following icons (from left to right):

- **Object Tree** opens the object tree window (see Object Tree Menu).
- **Host List** opens the host list window (see Host List).
- **Consoles** opens the consoles window (see Consoles (z/OS only)).
- **Add/Delete Working Sets** opens the working sets window (see Working Sets).
- **Active Objects List** opens the active objects list window (see Active Objects List).
- **Scheduler** launches the BrightStor Resource Manager Scheduler (see Scheduler).
- **Commands** opens the Operator Commands table.
- **Edit Member (Submit)** opens the Edit Member dialog (see The Member Editor (z/OS only)).
- **System Activity Log** opens the System Activity Log table (see the online help for details).
- **Minimize and Lock** minimizes BrightStor Resource Manager and locks it (see Scheduler Security (Locking the Application)).
- **Help** launches the online help.
- **Exit** closes BrightStor Resource Manager.

You can resize and reposition the toolbar as needed. Next time you start BrightStor Resource Manager, the toolbar appears with the same size and in the same location as it was when you quit the program.

Scheduler

Use the BrightStor Resource Manager Scheduler to activate scheduled objects and control scheduler functions.

To start the Scheduler:

✓ Click the Scheduler (clock) icon on the toolbar or select **Scheduler** from the **Vantage** menu.

To start the Scheduler each time BrightStor Resource Manager is started:

1. Select **Activation** from the **Tools** menu.

2. On the **General** tab, check the **Start Scheduler** box.

   Next time you start BrightStor Resource Manager, the Scheduler is started minimized and appears as a clock icon in the task bar.

Note: The Scheduler must be active for scheduled operations to be enabled.

Object Scheduling

You can schedule operations relating to any BrightStor Resource Manager view. Object scheduling is performed using the **View Definition** function (see View Definition).

Working with the Scheduler

The Scheduler activates schedules and controls their execution.

The BrightStor Resource Manager Scheduler window contains a scope pane that lists the various scheduled objects and a result pane that shows details of the object selected in the scope pane. Objects in the scope pane can be expanded to show their contents.

At the top of the object hierarchy is the object type (Trend or HTML). Below are the individual views (schedules) defined for each type. Below the views are the hosts on which the view has been defined.

- ✓ Click on an object to expand it and see the views defined under it.
- ✓ Click on a view to expand it and see the hosts on which it has been defined.

You can also use the **Expand** and **Collapse** icons in the toolbar or the corresponding functions in the right mouse menu to expand and collapse items in the scope pane.

If you select an object or a view in the scope pane, the result pane displays the information corresponding to your selection. If in the scope pane you select a view connected to a given host, the result pane shows detailed information about the selected view: its status and the relevant information according to its status (if the status is stopped, the date and time of the last and next updates; if the status is active, the last update and the number of processed records, and so on).

Right Mouse Menu

A status-sensitive right mouse menu is available in the scope pane for the following operations:

- Activate or disable a schedule
- Stop a running process
- Expand and collapse the tree in the scope pane
- Launch the online help
- Close the Scheduler

The menu is also available from the result pane when the pane shows a list of objects.

Scheduler Security (Locking the Application)

To prevent security breaches while the Scheduler is active, you can use the minimize and lock function available from the main toolbar.

✓ Click the **Lock** icon (Minimize and Lock) on the main toolbar.

BrightStor Resource Manager is minimized to an icon in the Windows task bar. BrightStor Resource Manager is running but it is locked. To unlock the application, you must enter your user name and password:

1. Click the BrightStor Resource Manager icon on the Windows task bar.
2. When the login dialog appears, enter your user name and password.

The Member Editor (z/OS only)

Use the BrightStor Resource Manager PDS member editor to edit configuration files, JCL templates, or any other text files. The generic editor supports standard text editing functions such as cut, copy, and paste, save and save as. For specialized objects, such as JCL templates, the functions of the editor are extended to support specialized operations such as variable substitution.

Opening a Member

The Open Member dialog lets you retrieve the desired member from a PDS.

1. Select **Edit Member** from the **z/OS** menu.
2. When the New PDS Member dialog appears, select **Open** from the **File** menu or click the **Open** icon.

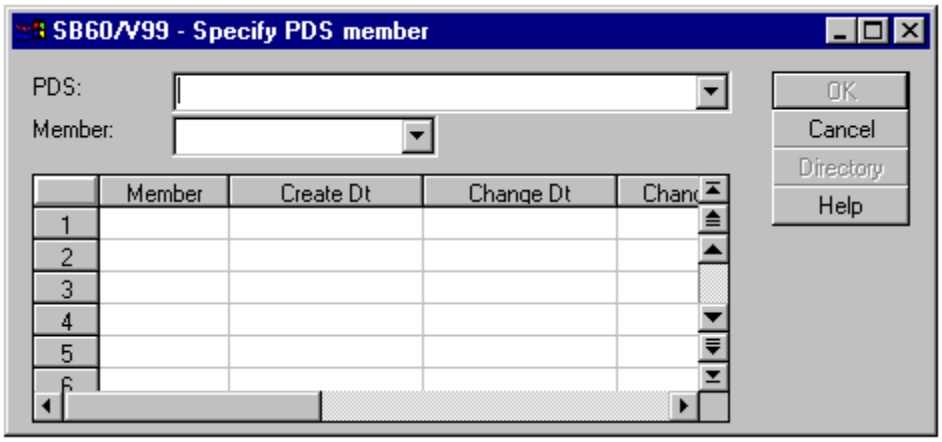

3. When the Specify PDS Member dialog appears, enter the name of the desired PDS and member and click **OK**.

   You can obtain a list of all the members in a selected PDS by clicking the **Directory** button. The system displays a list of members in the selected directory, including the date and time of the last modification, member size, and user ID.

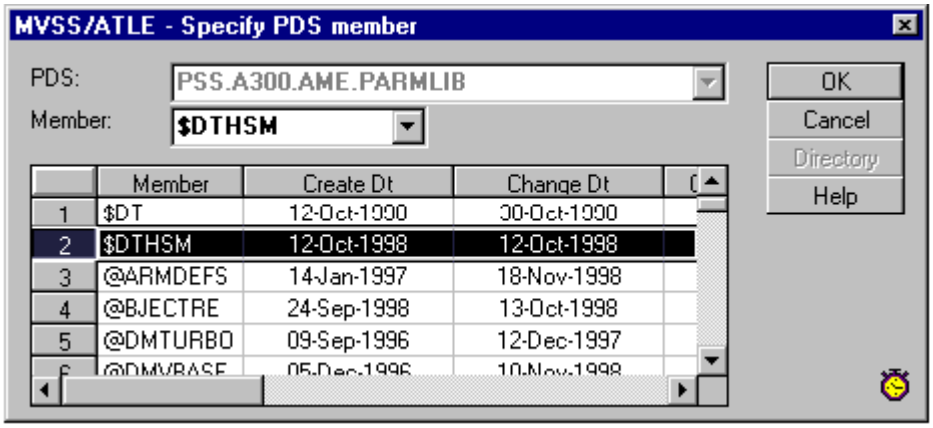

If you do not supply a member name when you click **OK**, the system attempts to open the data set specified in the **Member** text box as a PS file.

After you click **OK**, the system brings the selected file or member into the Edit Member window.

Editing a Member

The Edit Member window has its own menu, which offers a set of generic editing tools.

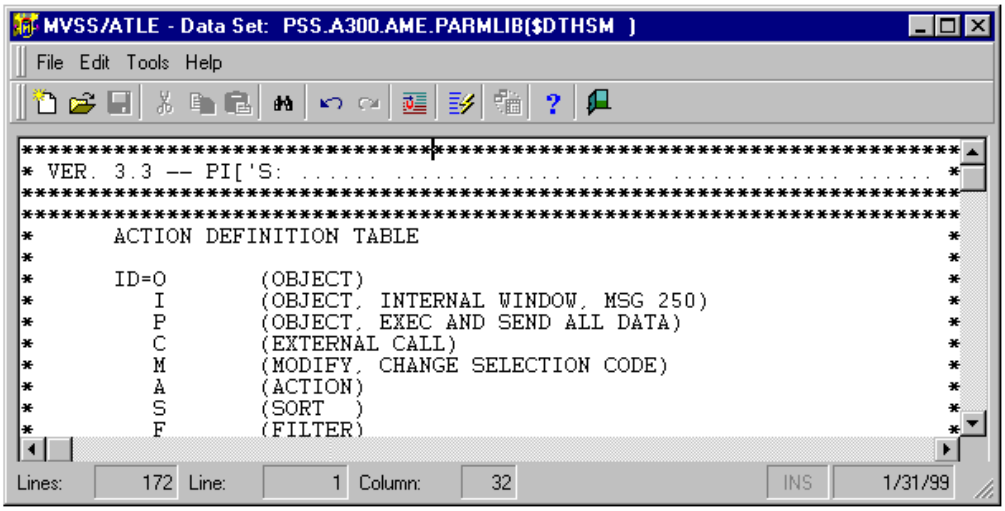

The editor menu has the following structure:

File	Edit	Tools
New	Undo	Variables
Open	Redo	Submit
Save	Cut	Show substitution
Save As	Copy	
Exit	Paste	
	Find	
	Replace	

The most frequently used functions are also available from the toolbar that appears below the menu. Placing the cursor above an icon in the toolbar for three seconds displays a balloon help label that describes the action of the icon.

The editor provides access to the following set of standard text and file manipulation commands:

- ✓ Use the **Cut**, **Copy**, and **Paste** commands in the **Edit** menu to perform standard text editing functions.

- ✓ Use the **Find**, **Find Next**, and **Replace** commands in the **Edit** menu to search for desired text in the member and to perform global text replacements within the member.
- ✓ Use the **Undo** command in the **Edit** menu to reverse the last editing operation you performed.
- ✓ If you want to open a new member, use the **Open** command in the **File** menu to bring up the Open Member dialog.
- ✓ If you modified a member, use the **Save** command in the **File** menu to save the member under its existing name, or **Save As** to save it under a new name.
- ✓ Use the **Close** command in the **File** menu to exit the editor without saving the current member.
- ✓ Use the **Submit** command in the **File** menu to submit a JCL for immediate or deferred execution.
- ✓ Use the **Variables** command in the **Edit** menu to invoke the JCL Variables list, a typing aid that simplifies the insertion of variables into JCL templates.
- ✓ Use the **Substitute** command to substitute the variables in the templates with values extracted for selected objects.

Find

The **Find** command lets you search for a text string in a member displayed in the editor.

1. Click on the **Find** icon, or select **Find** from the **Edit** menu to activate the **Find** command.
2. When the Find dialog appears, enter the string you want to search for in the **Find What** box.
3. To make the search case sensitive or limit it to occurrences of the search string in whole words only, click the **Case Sensitive** and **Whole Words Only** check boxes.
4. Select the **Up** or **Down** radio boxes to determine the direction in which the search will proceed in the member.
5. Click the **Current Location** or **Beginning of Module** buttons to determine whether the search begins at the location of the cursor or at the beginning of the file.

After you click **OK**, the cursor skips to the first occurrence of the word, and highlights it. You can repeat the search for further occurrences of the string by selecting **Find Next**.

Replace

The **Replace** command lets you search for a text string in a member displayed in the editor and replace it with another string.

1. Select **Replace** from the **Edit** menu.

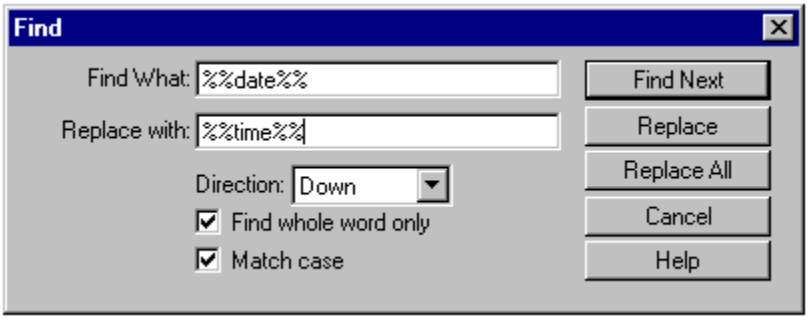

2. When the Find dialog appears, enter the string you want to replace in the **Find What** text box and the string you want to replace it with in the **Replace With** text box.
3. To make the search case sensitive or limit it to occurrences of the search string in whole words only, click the **Case Sensitive** and **Whole Words Only** check boxes.
4. Select the **Up** or **Down** radio boxes to determine the direction in which the search will proceed in the member.
5. Click the **Current Location** or **Beginning of Module** radio buttons to determine whether the search begins at the location of the cursor or at the beginning of the file.

After you click **OK**, occurrences of the string in the **Find What** text box are replaced with the string in the **Replace With** box.

Inserting Variables

To insert a JCL variable:

1. Click the **Variables** icon or select **Variables** from the **Tools** menu

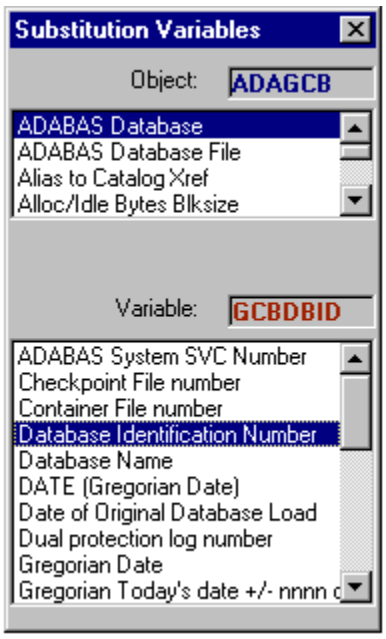

   The Variables window appears. Variables in the window are arranged by object type.

2. Click one of the object types in the top half of the window to obtain a listing of the variables available for that object in the bottom part of the window.

3. Double-click the desired variable.

   The selected variable, embedded within double % signs, is transferred to the editor at the location of the cursor.

By including a REPEAT statement (or a pair of REPEAT/REPEATEND statements) in the JCL template, you can perform the same action on more than one record (row) within an object table.

REPEAT causes the substitution mechanism to duplicate the last section of the template for each object on which you want to perform the operation. (REPEAT duplicates the last portion of the template as many times as there are selected records in the object table; variables that appear before the REPEAT statement are substituted only for the first object selected.)

The pair REPEAT/REPEATEND causes the substitution mechanism to duplicate the section of the template located inside the pair as many times as there are selected records. Variables that appear outside the REPEAT/REPEATEND pair are substituted only for the first record selected.

The syntax of the REPEAT statement is:

```
//*REPEAT*
```

In the following example, the REPEAT statement appears before the %%VOL%% variable.

DFSMSdss utility to defrag a volume:
(In JCLLIB member=SAMPJCL4)

```
//SSM%JOB JOB (.... your job statement ..............
//*************************************************************
//* ** EXAMPLE OF HOW TO RUN DEFRAG AGAINST VOLUME    ********
//* **                                                ********
//* ** NOTE!VOLUMES HAVING PROCLIBS AND LINK LIST DATA SETS  **
//* **       SHOULD NOT BE CANDIDATES FOR DEFRAG, OR USE    **
//* **       EXCLUDE PARAMETER FOR THOSE DATA SETS          **
//*************************************************************
//DEFRAG   EXEC PGM=ADRDSSU
//SYSPRINT DD SYSOUT=*
//SYSIN    DD *
//*REPEAT*
DEFRAG     DYNAM(%%VOL%%)
```

If three volumes are selected, say SYSC01, SYSC02, and SYSC03, the variable %%VOL%% is substituted three times, once for each volume, the DEFRAG statement is repeated three times and the result appended to the substituted JCL. If you select only one record in the object, the REPEAT statement has no effect.

The syntax of the REPEAT/REPEATEND statement is:

```
//*REPEAT*
//%%YOUR_VARIABLE%%
//*REPEATEND*
```

In the following example, the REPEAT/REPEATEND pair appears twice, first around the %%VOL%% and %%DSN%% variables where it causes all the selected data sets to be deleted, and again following the **Mail** command, where it sends a message about each deleted data set to the designated BrightStor Resource Manager user.

```
//JOBCARD.....................................
//IEFBR14   EXEC PGM=IEFBR14
//*REPEAT*
//%%VOL%%   DD DSN=%%DSN%%,DISP=(OLD,DELETE)
//*REPEATEND*
/*
//MAILBOX   EXEC PGM=VANSENDM
//STEPLIB   DD DISP=SHR,DSN=VANTAGE.LOADLIB
//PARMS     DD DISP=SHR,DSN=VANTAGE.PARMLIB
//VANMSG    DD *
MAIL_TO_USER=SSSAMY,SEV=I,MSG=HAVE DELETED TEST DATA SETS
//*REPEAT*
DATASET NAME %%DSN%% ON VOLUME %%VOL%%
//*REPEATEND*
/*
```

Substituting Data in the Variables

You can substitute real values for the template variables by using the BrightStor Resource Manager selection tools to select the objects that contain the real data. Suppose the action you are defining is a backup and the substitution variable is DSN (data set name). To generate the JCL code that executes the backup, substitute the names of all the data sets you want to back up for the %%DSN%% variable in the template.

1. Open the Data Sets table and use the appropriate sorting and filtering commands to narrow the selection of data sets that appear in the table.
2. Select the data sets you want included in the backup action.
3. Click in the first column of one of the selected lines, then drag the item into the Edit Member window and drop it in (release the mouse button).

   When you drop the selected lines into the Edit Member window, the list of selected objects is sent to the host, where a substituted member is created as a memory table. When the PC receives the message from the host that the substituted member has been created, a new button, **Refresh**, appears in the toolbar of the Edit Member window.
4. To view the substituted member, click the **Refresh** button.

You can now review the substituted member, save it, or submit it for execution. You can dismiss the window containing the substituted member by clicking its control box, or reduce it to an icon by clicking the **Minimize** button.

Saving

To save the edited member:

✓ Select **Save** from the **File** menu.

To save an edited member under a different name for later reuse:

1. Select **Save As** from the **File** menu.
2. When the Save Member dialog appears (identical with the Open dialog), specify file and member names.
3. Click **OK**.

Using the Submit Command

The **Submit** command lets you submit a JCL for execution.

1. Select **Submit** from the **Tools** menu.

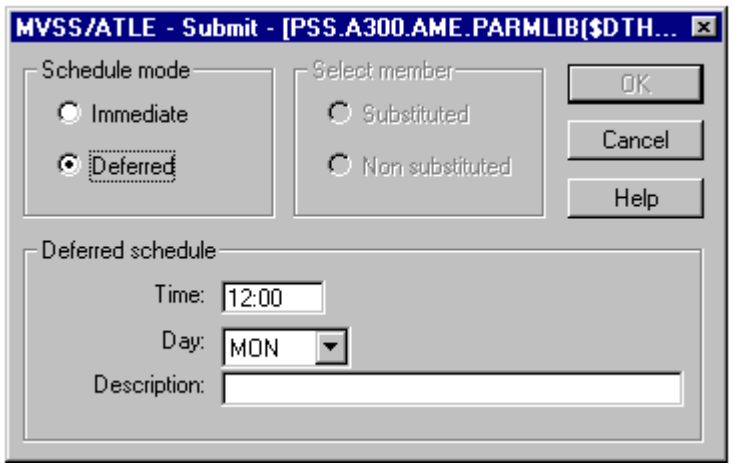

2. When the Submit dialog appears, click the **Immediate** radio button to submit the job immediately, or the **Deferred** button to schedule it for later execution.
3. If you clicked the **Deferred** button, assign execution day and time and provide an optional description of the job in the appropriate fields of the **Deferred Schedule** area of the screen.
4. In the **Select Member** section, click the **Substituted** button to submit the substituted member (which is not shown in the editor's window), or the **Non substituted** button to submit the original member shown in the editor's window.
5. Click **OK** to complete the submission.

Audit and System Scripts (z/OS only)

BrightStor Resource Manager generates audit reports about the discrepancies between entries in files and catalogs. The following audit reports are available:

- DFSMShsm BCDS Audit Report
- DFSMShsm MCDS Audit Report
- DFSMShsm OCDS Audit Report
- DFSMShsm MCB/MCC Audit Report
- Tape Management Systems
- Uncataloged Data Sets
- Audit of Cataloged Data Sets

BrightStor Resource Manager also provides the following system scripts:

- VSAM VVDS Extract
- ACS/LSM Volume Report
- TMC DSNB Performance
- TMS/Silos Scratch Tapes Mismatch
- Group Space on Primary Storage
- Duplicate Data Sets for System
- Multi Volume Data Sets
- Disk Checkpoint Data Sets

Wizards help you define and activate scripts for each of the operations listed above. (The operation of the wizards is identical for all scripts, with one additional panel used for Tape Management System audits.)

Defining Audit Scripts

Audit operations are carried out by automated scripts. BrightStor Resource Manager provides a wizard to help you define and activate scripts for each of the audits listed above. (The operation of the wizards is identical for all audit scripts.)

To define a new audit script or to modify an existing one:

1. From the main BrightStor Resource Manager **z/OS** menu, select **Object Scripting**, then **Audit Scripts**, then one of the audit scripts in the list.
2. When the Wizard panel appears, enter the system ID of the host in the appropriate field, and click **Next**.
3. Select a day from the **Day** drop-down box and enter hour and minute values in the **Time** boxes to specify on which day and at what hour the script should be enabled on the host. Click **Add to list**.

   The day and time appear in the timing list. You can add as many execution days/times to the list as needed. You can remove an entry from the list by selecting it and clicking **Remove**.
4. Check the **Execute the script every time...** button to execute the script each time it is activated. If this box is checked, the **Save, Activate, and Execute** button is disabled. If you leave the box unchecked, all buttons are enabled.
5. Click **Save but do not Activate** to save the script without activating it (the script is not enabled on the host).
6. Click **Save and Activate** to enable the script when the next active period defined in the day and time fields is reached.

7. Click **Save, Activate, and Execute** to fire the script immediately after the wizard completes execution. Click **Finish**.

Tape Management System Audits

The Tape Management System Audits wizard lets you define several audit types per tape management system. To define a new audit script or to modify an existing one:

1. From the main BrightStor Resource Manager **z/OS** menu, select **Object Scripting**, then **Audit Scripts**, then **Tape Management Systems**.
2. When the Wizard panel appears, enter the system ID of the host in the appropriate field, and click **Next**.
3. When the Tape Management System panel appears, select one of the tape management systems by clicking its button and select one or more audit types. Click **Next**.
4. Select a day from the **Day** drop-down box and enter hour and minute values in the **Time** boxes to specify on which day and at what hour the script should be enabled on the host. Click **Add to list**.

   The day and time appear in the timing list. You can add as many execution days/times to the list as needed. You can remove an entry from the list by selecting it and clicking **Remove**.
5. Check the **Execute the script every time...** button to execute the script each time it is activated. If this box is checked, the **Save, Activate, and Execute** button is disabled. If you leave the box unchecked, all buttons are enabled.
6. Click **Save but do not Activate** to save the script without activating it (the script is not enabled on the host).
7. Click **Save and Activate** to enable the script when the next active period defined in the day and time fields is reached.
8. Click **Save, Activate, and Execute** to fire the script immediately after the wizard completes execution. Click **Finish**.

Defining System Scripts

BrightStor Resource Manager provides system scripts to perform the following operations:

- VSAM VVDS Extract
- Uncataloged Data Sets
- Audit of Cataloged Data Sets
- ACS/LSM Volume Report
- TMC DSNB Performance

- TMS/Silos Scratch Tapes Mismatch
- Group Space on Primary Storage
- Duplicate Data Sets for System
- Multi-Volume Data Sets for System
- Disk Checkpoint Data Sets

A wizard helps you define and activate scripts for each of the operations listed above. (The operation of the wizards is identical for all scripts.)

Note: You can launch the wizard in two ways: from the **z/OS** branch of the main BrightStor Resource Manager menu and from the **Actions** menu of individual tables. For example, to launch the wizard that configures the automation script for duplicate data sets, you can follow the path **Object Scripting, System Scripts**, **Duplicate Data Sets** from the **z/OS** menu, or select **Configure Script** from the **Actions** menu of the Duplicate Data Sets table.

To define a new system script or to modify an existing one:

1. Launch the wizard using one of the two methods described above.
2. When the Wizard panel appears, enter the system ID of the host in the appropriate field, and click **Next**.
3. Select a day from the **Day** drop-down box and enter hour and minute values in the **Time** boxes to specify on which day and at what hour the script should be enabled on the host. Click **Add to list**.

   The day and time appear in the timing list. You can add as many execution days/times to the list as needed. You can remove an entry from the list by selecting it and clicking **Remove**.
4. Check the **Execute the script every time...** button to execute the script each time it is activated. If this box is checked, the **Save, Activate, and Execute** button is disabled. If you leave the box unchecked, all buttons are enabled.
5. Click **Save but do not Activate** to save the script without activating it (the script is not enabled on the host).
6. Click **Save and Activate** to enable the script when the next active period defined in the day and time fields is reached.
7. Click **Save, Activate, and Execute** to fire the script immediately after the wizard completes execution. Click **Finish**.

Note: Audit scripts also can be accessed for editing and execution from the various Tape and Robotics audit tables.

TMS/Silos Scratch Tape Mismatch

Use the TMS/Silos Scratch Tape Mismatch wizard to define several audit types for each tape management system. To define a new system script or to modify an existing one:

1. From the main BrightStor Resource Manager **z/OS** menu, select **Object Scripting**, then **System Scripts**, then **TMS/Silos Scratch Tape Mismatch**.
2. When the Wizard panel appears, enter the system ID of the host in the appropriate field, and click **Next**.
3. When the Tape Management System panel appears, select one of the tape management systems by clicking its button and select one or more audit types. Click **Next**.
4. Select a day from the **Day** drop-down box and enter hour and minute values in the **Time** boxes to specify on which day and at what hour the script should be enabled on the host. Click **Add to list**.

   The day and time appear in the timing list. You can add as many execution days/times to the list as needed. You can remove an entry from the list by selecting it and clicking **Remove**.
5. Check the **Execute the script every time...** button to execute the script each time it is activated. If this box is checked, the **Save, Activate, and Execute** button is disabled. If you leave the box unchecked, all buttons are enabled.
6. Click **Save but do not Activate** to save the script without activating it (the script is not enabled on the host).
7. Click **Save and Activate** to enable the script when the next active period defined in the day and time fields is reached.
8. Click **Save, Activate, and Execute** to fire the script immediately after the wizard completes execution. Click **Finish**.

Logging and Capturing Data (z/OS only)

Note: Trend logs for volumes, data set groups, external groups, and storage groups created with versions of BrightStor Resource Manager earlier than 4.0 are not compatible with the new trend and reporting system included in the current release. Information on the log conversion utility is provided to convert the old trend logs to the new format and transmit records to the Windows Client for trend analysis.

You can use two methods for logging the contents of the current object to a sequential data set on the host: **Logging** and **Capture**. Logging is intended to be used by storage administrators for general storage management purposes. Capturing is intended for individual users to take the snapshots of object data on an *ad hoc* basis.

For each object, BrightStor Resource Manager maintains one log file and one capture file. Data from successive logging operations is appended to the same log file; similarly, all captured data is stored in the same capture file. The object name is used as part of the dsname for the sequential data set.

To log the data of BrightStor Resource Manager objects, you can define and manage logging scripts.

To capture the data of a specific object, use the **Export** function described in Graph Definition.

The logged and captured object data can be viewed by selecting **Mode** from the object table's **View** menu and choosing the appropriate data collection mode. You must specify the range of dates and times for which the log or capture data is to be displayed (as described in Mode (z/OS)).

Defining Log Operations

To log the data you want to use in trend analysis, you must tell the system what to log and when; for this purpose BrightStor Resource Manager provides a Log Operations Wizard. Logging is available only for objects listed in the Log Operations Wizard.

To define a log operation:

1. Select **Object Trend Logging** from the **z/OS** menu (in the main toolbar) or from the right mouse button menu of the object tree.

   Alternatively, right-click on a desired object in the object tree and select **Object Logging** from the Wizard pop-up menu.

2. When the Log Operations wizard's opening screen appears, click **Next**.

3. Select the desired objects from the tree of objects for which logging is available and click **Next**.

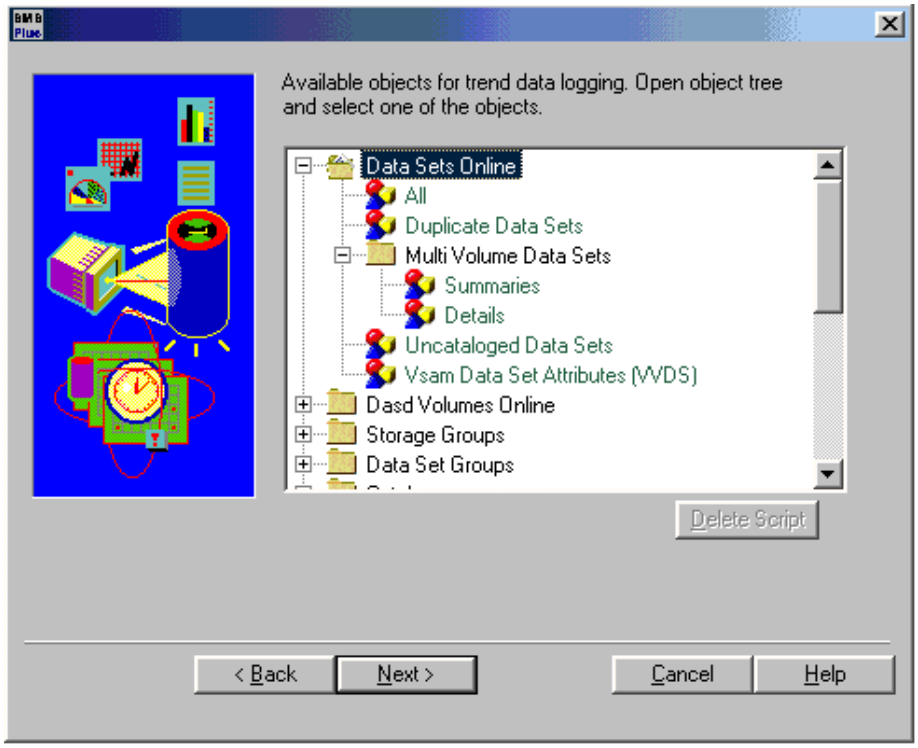

4. Select the timing for the logging operation.

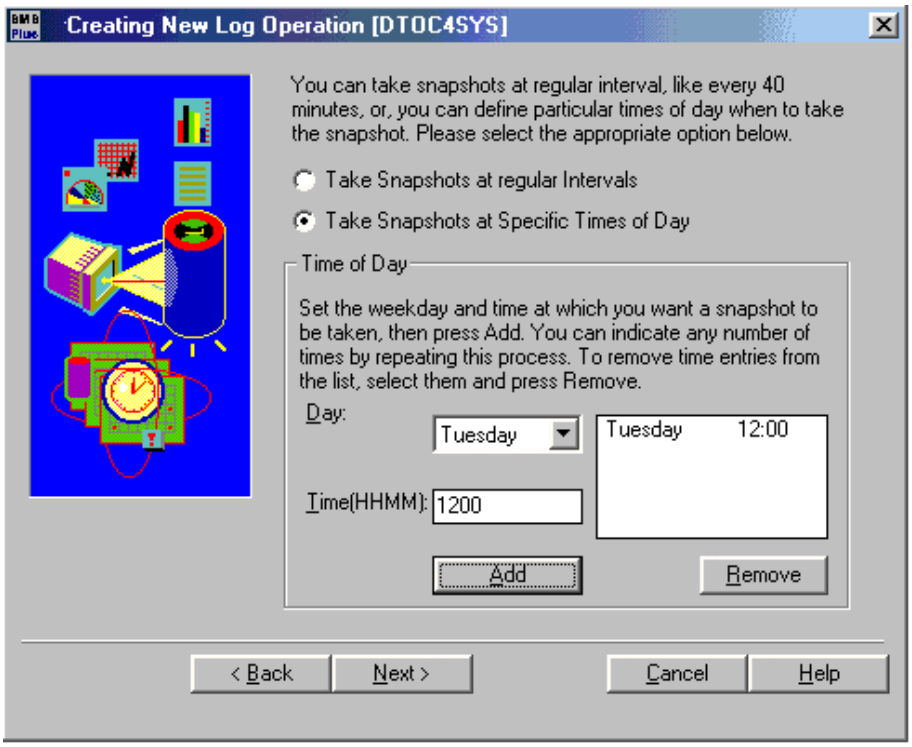

You have two options:

- To request logging at regular time intervals, click the **Take Snapshots at Regular Intervals** button. BrightStor Resource Manager then prompts for the time interval value.
- To request logging at specific times, click the **Take Snapshots at Specific Times of Day** button. Specify the list of times and days of the week when object snapshots are to be taken. You can modify the list using the **Add** and **Remove** buttons.

5. Click **Next**.
6. Define the filter to specify the instances of the object you want logged. (This step is optional; if skipped, the snapshots of all the instances of the selected object are taken.) The **Define Filter** button invokes the standard Filter dialog (see Filtering for information on how to define filters). Enter a title and an optional description for the script. Enter the system ID of the host in the appropriate field. Click **Next**.
7. Specify when to create a new log file and the maximum number of logs per file. Check the appropriate box(es) to log data, object count, or both. Click **Next**.

8. Specify whether to collect *ad hoc* data before taking the snapshot of the object (click the **Base Snapshot on Real Time** radio button), or use data acquired during the last scanning interval (click on the **Base Snapshot on Last Interval** radio button).

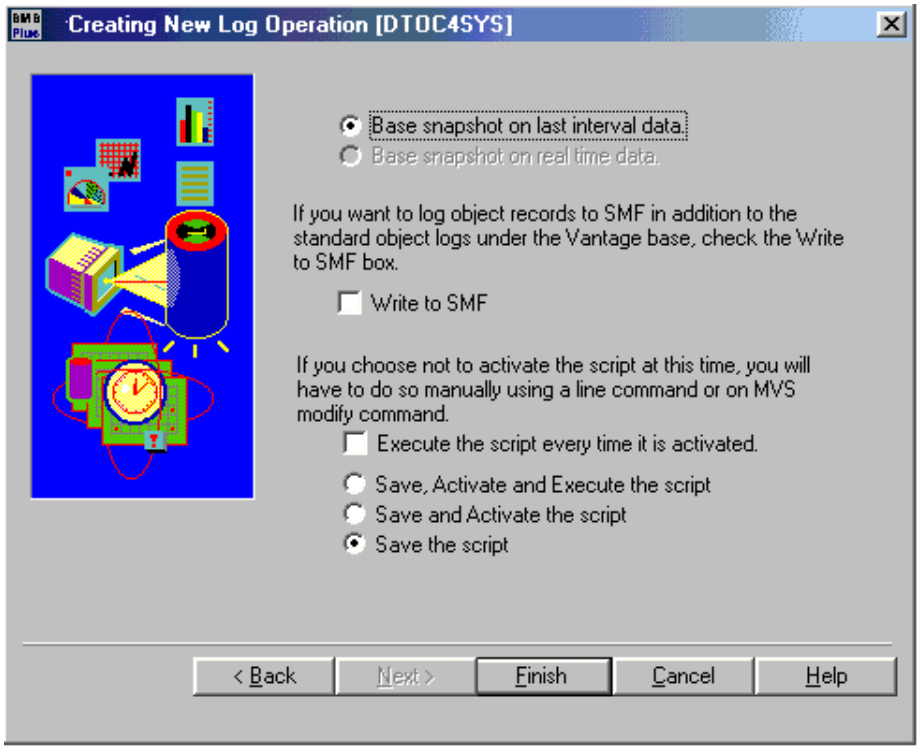

9. To log object records to SMF in addition to the BrightStor Resource Manager object logs, check the **Write to SMF** box.

   Note: If you select to write records to SMF, review and customize the BrightStor Resource Manager system parameter LOGSMFRT.

10. Check the **Execute the script every time...** button if you want the script to be executed each time it is activated. (If this box is checked, the first button, **Save, Activate, and Execute** is disabled. If you leave the box unchecked, all buttons are enabled.)

11. Click **Save, Activate, and Execute the script** to activate and execute the script as soon as it is saved, or **Save and Activate the script** to save and activate the script without executing it, or **Save the script** to save the script without activating it. If you choose not to activate the script, you must activate and execute the script manually, using either a line command or an MVS modify command.

12. Click **Finish**.

The wizard displays the message that the specified logging script data was saved.

Managing Log Operations

You can start, stop, and modify the defined logging operations. To manage logging operations:

✓ Select **Modify Existing Log** from the **Object Trend Logging** branch of the **z/OS** menu.

BrightStor Resource Manager displays the Log Operations wizard. Make the necessary corrections and click **Apply** button. Refer to the previous section for the description of Log Operations wizard's screens.

System Parameters (z/OS only)

To reach the Parameters table:

✓ Select **System Parameters** from the **z/OS** menu.

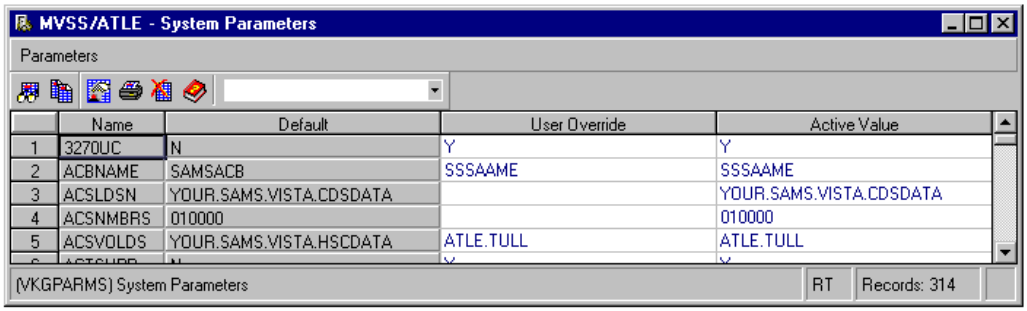

PARMDEFS and When the Parameters table appears, the mainframe creates a system parameters table by merging the memory table with the information from VKGPARMS and sends the complete table to the PC.

For each parameter, the Parameters table displays a row of information. The table contains the following fields:

Field	Description
Name	Parameter name.
Default	Default, pre-installed value in member PARMDEFS.
Active	Parameter value in memory table.
Vkgparms	Parameter value in member VKGPARMS (if any).

The parameters displayed in the table are sorted alphabetically by name. You control the table display in the following ways:

To show parameters related to one of the system objects:

✓ Select that object from the **Objects** list box and click **Execute**.

When the parameters of the selected object appear in the table, you can click anywhere in a row to see a brief explanation of that parameter in the status line, at the bottom of the table.

You can edit the values of system parameters stored in the VKGPARMS member and in the memory table as follows:

1. When the parameter you want to edit is visible on the screen, move the cursor to the desired parameter row in the **Vkgparms** or **Active** column.

   You can modify the values in these columns.

2. Press **F1** when the cursor is on a parameter line to display online help for the parameter.

To delete a parameter from VKGPARMS (reverting its value to the default):

✓ Place the cursor in the desired cell and clear its value by pressing the **Del** key.

To save all modifications in the VKGPARMS member and to update the memory table:

✓ Click the **Save** button. **Save** sends all the lines that have changed from the Windows Client to the host. If any parameter members have been modified, the host updates the system parameters memory table and rewrites the VKGPARMS member into PARMLIB.

   Additionally, if you select a new object or if you choose to exit the table, the system asks you whether or not to save the changes you have made.

To reload the values from VKGPARMS into the memory table:

1. Select **Commands!** from the BrightStor Resource Manager main menu.
2. When the Operator Commands dialog box appears, select REFRESH,VKGPARMS.

To print the list of parameters displayed in the table:

✓ Click the **Print** button.

Operator Commands (z/OS only)

You can use operator commands on selected objects through the **Commands!** menu. These commands cover all object-related system commands in the BrightStor Resource Manager subsystem.

To execute a command:

1. Select **Commands!** from the given object's menu. BrightStor Resource Manager displays the Commands window. If you select **Commands!** from the main menu, the window lists all available commands. If you select **Commands!** from an object table (for example, storage groups or data set groups), the window lists the commands available for that object.

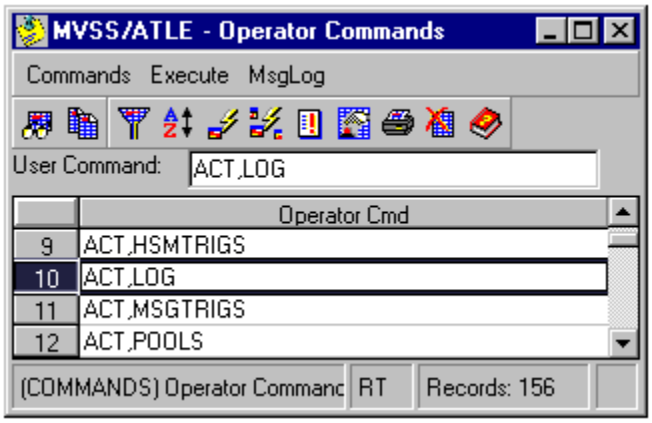

2. Select a command from the command list. BrightStor Resource Manager moves the selected command to the **User Command** box below the toolbar.
3. Modify the command as necessary.
4. Select **Execute** from the menu bar or click the **Execute** button. BrightStor Resource Manager executes the command. To execute the command on all selected hosts simultaneously, click the **Multi Execute** button and confirm the action.

Note: BrightStor Resource Manager does not inform you whether the command was executed successfully. The only message that appears is **Command submitted.** However, you can verify the result of the command by checking the BrightStor Resource Manager Log, which you can reach directly from this window with the **MsgLog** button or menu selection.

5. Select another command to execute or click the **Close** button to close the Commands window.

Note: When you issue a command from one of the BrightStor Resource Manager online interfaces, the first two words and the comma are omitted.

The standard BrightStor Resource Manager table functions, **Save Save As**, **Print**, and **Export** are available from the Operator Commands table.

Available Operator Commands

The complete list of operator commands is included in the chapter Operator Commands in the *BrightStor Resource Manager Messages and Reference Manual.*

Note that only commands that are available to Windows users are displayed when you select **Commands!** from the main toolbar. Several commands are available only from the MVS operator console and do not appear under the Windows interface.

Object Properties

You can reach the Object Properties dialog from the right mouse menu. The dialog provides access to the following functions:

- Color definition
- Scale base selection

Color Definition

Use the Color Definition dialog to control the color coding of individual fields in a view. For numeric fields, you can assign custom colors to values that are displayed as bars. For all fields you can assign a color to a condition and display that field in the selected color if the condition is met.

1. Select a view, then choose **Object properties** from the right mouse button menu. When the Properties dialog appears, click the **Colors** tab.

The Object Properties dialog shows a list of all the fields of the selected view.

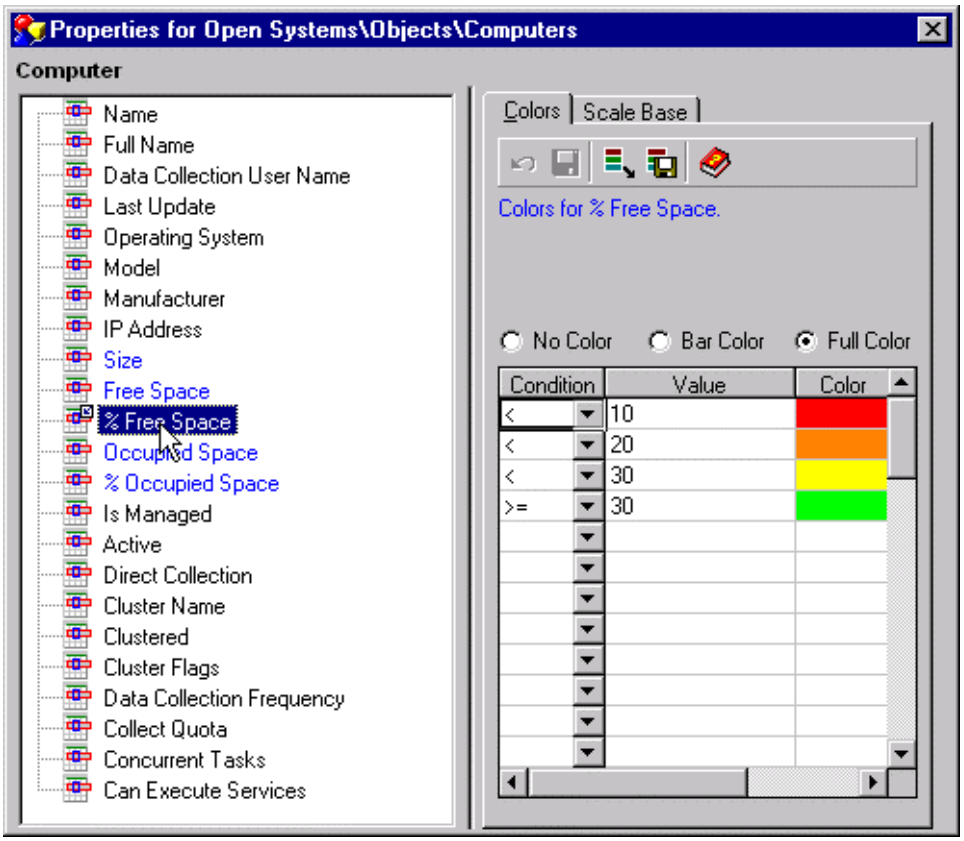

2. Select one of the fields.

   In the right panel of the dialog, the default **No Color** button is selected.

3. Click the **Bar Color** or **Full Color** button, to define the color of bars in numeric fields or the color in which the entire field is rendered.

   The display expands into a table with columns for conditions and matching colors. Conditions consist of a logical operator and a value.

4. Define a condition by selecting a logical operator from the **Condition** column and entering a value in the **Value** column.

5. Assign a color to the condition by clicking the cursor in the **Color** column.

6. When the Color dialog appears, click on the desired color and click **OK**.

7. To define and apply a color other than the ones that appear in the list of **Basic colors**:

   a. Click the **Define Custom Colors** button.

      The dialog expands to the right to let you define custom colors and assign them to one of the 16 squares in the **Custom colors** group.

b. Select one of the **Custom colors** squares.

c. Set the **Luminosity** slide bar to the desired level.

240, at the top of the scale, corresponds to white; 0, at the bottom of the scale, corresponds to black. Note the values changing in the **Lum**(inosity) and **Red**, **Green**, and **Blue** fields as you move the slide bar.

d. Move the cursor on the color palette over the desired color.

Note the values changing in the **Hue** and **Sat**(uration) fields. Values in the **Red**, **Green**, and **Blue** fields change within the limits of the selected luminosity.

Your color selection appears in the **Color|Solid** box.

e. Click the **Add to Custom Colors** button when you are satisfied with the selected color.

The new color appears in the **Custom colors** group on the left side of the dialog.

8. Click **Save** in the Color Definition dialog to save the definition for the selected field.

You can apply color definitions you created for one field to other fields. To do so:

1. Complete a color definition for any field and click Save as Default.
2. Select another field and click **Use Default**. The definition that has been saved as default is applied to the new field.

Scale Base Selection

Use the Scale Base selection tab to specify the unit in which data is displayed in numeric fields.

1. Select a view, then choose **Object properties** from the right mouse button menu. When the Properties dialog appears, click the Scale Base tab.

   The Scale Base tab shows a list of all the fields of the selected view.

2. Select one of the numeric fields.
3. From the **Scale Base** drop-down list select one of the following units of measure for storage entities:

   Decimal units such as K (1,000), M (1,000,000), etc.

   Binary units such as KB (1,024), MB (1,048,576), etc.

   Individual units (not modified by any multiplier).

   The original units in which the value is reported by the host.

Note: If you have been collecting trend data using a given unit of measure, changing the scale base can affect the appearance of the trend report and previously collected data may not appear on the charts.

Activating Components

Use this dialog to specify the BrightStor Resource Manager features and components you want active the next time the system is started.

Information in the dialog is organized into several tabs:

- Use the General tab to specify the windows you want open when BrightStor Resource Manager is started. You can choose to have the Object Tree, Host List, and Views Status windows open at start-up. If you check the **Run in Demo mode** box, BrightStor Resource Manager starts without connecting to a host. If any tables have been saved with the **Save Data for Demo** function (see Save Data for Demo), these tables can be accessed in demo mode.

  If you check the **Automatic Table Column Adjustment** box, the BrightStor Resource Manager automatically adjust the width of all the columns in the table to the width of the data.

  If you check the **Automatic Table Column Adjustment box**, the width of all the columns in all the tables is automatically adjusted to the width of the data.

  If you check the **Start Scheduler** box, the BrightStor Resource Manager Scheduler is started each time you start BrightStor Resource Manager.

  If you check the **Inhibit scheduler from disconnecting hosts** box, the Scheduler does not break the client connection to the host after completing a scheduled task, even if no other objects are active on that host.

  Use the **Data Selection Method in Tables** radio buttons to determine the method of cell and row selection in BrightStor Resource Manager tables.

  - In **Block Mode**, you can select individual lines and cells or drag the mouse to select any rectangular block of cells. By holding down the **Shift** key you can extend the size of a contiguous rectangle or define the rectangle by clicking on two opposite corners. By holding down the **Control** key, you can select several non-continuous blocks. Note that you cannot unselect individual blocks.
  - In **Line Mode**, you can click the mouse to select individual rows (lines), **Shift-**click to select a group of contiguous rows, or **Ctrl**-click to select non-contiguous rows.
- Use the Hosts tab to specify the host connections you want to be active when BrightStor Resource Manager is started.

- Use the User Views tab to specify the user views you want to be active when BrightStor Resource Manager is started. Available views are listed in the top pane; selected views in the bottom pane. You can move views between the two panes by selecting them and clicking the **up** and **down arrows**.
- Use the Working Sets tab to specify the working sets you want to be active when BrightStor Resource Manager is started. Check the box next to each working set you want to be active.
- Use the Consoles tab to specify the consoles you want to be active when BrightStor Resource Manager is started. Check the box next to each console you want to be active.

After you make your selections on all the tabs, click **OK**. Your choices become active the next time the system is started.

Consoles (z/OS only)

Consoles are collections of columns of information from one or more hosts. For each console you specify the hosts (a subset of the active sessions) and columns you want to display.

When more than one communication session is in progress, you can choose to display the consolidated information obtained from different sessions. For example, if you have active BrightStor Resource Manager sessions running under four different systems, and you are interested in monitoring hardware errors on all four, you can define a console that shows the number of hardware errors on all the active systems.

A console shows a subset of all active sessions and columns of information that have been defined. For each console you define, you specify the systems you want listed in the console and the columns of information to be included.

For each active session, the console provides at-a-glance information that is important to the operation of your site. In the console table, each system is listed as a row. The columns contain information you have decided to include in the display. The information in each cell can be a statistic such as the number or percentage of volumes that exceed a certain threshold, or it can represent a complex query performed on data collected from several system objects. In all cases, the value that appears in the columns represents a calculated value. The BrightStor Resource Manager console definition mechanism provides the tools you need to specify the objects and the range of data you want collected and to define the queries you want to perform on the data in order to produce the value that appears in each column.

A few column definitions are supplied with BrightStor Resource Manager. You can modify these definitions and add your own.

The BrightStor Resource Manager Consoles window provides the following services:

- Lets you define new consoles and modify existing ones.
- Shows the list of defined consoles and lets you select one of them.
- Provides a wizard for the definition of new columns.

A right mouse button menu provides access to the following functions:

- Expand and collapse the console and column trees
- Find
- Add a new console
- Add a new column
- Edit a console or column
- Delete a console or column

Defining Consoles

To add a new console definition to the list in the BrightStor Resource Manager Consoles window:

1. Click on the **Consoles** icon, then right-click the mouse button and select **New** from the right mouse button menu.

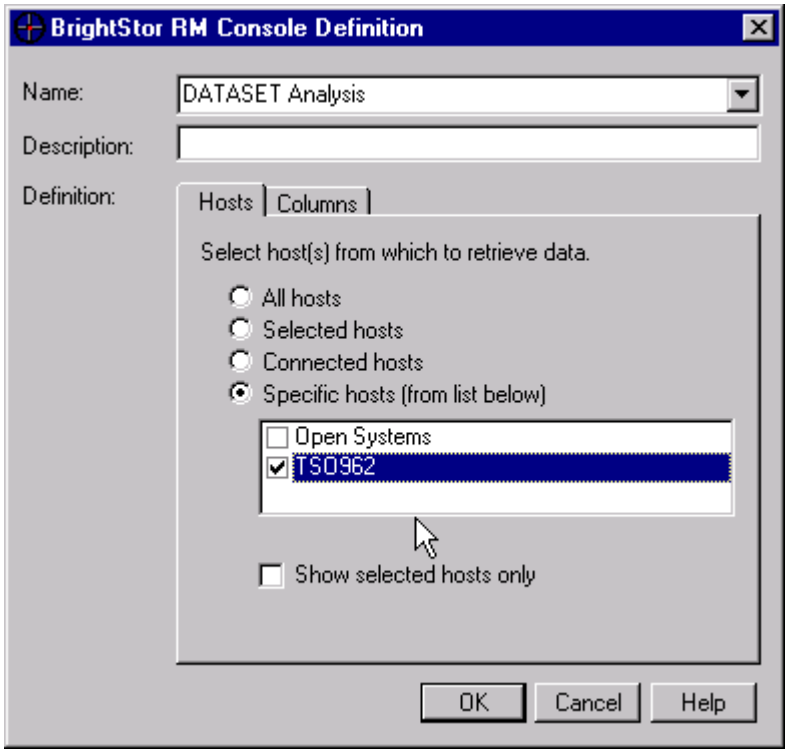

2. On the Hosts panel of the BrightStor Resource Manager Console Definition dialog, enter a name and description for the new console in the appropriate text boxes.

3. Select the hosts you want to include in the console view by checking their boxes. The name of each system that appears in the list is concatenated from the system's ID and subsystem name.

   You can click the **All hosts** button to select all defined hosts. If you click the **Connected hosts** button, data is automatically collected from all connected hosts and included in the console.

   Note: The sysid and the subsystem name are based on the system ID and subsystem name parameters that were entered using the Host Definition dialog. Make sure that the system ID and subsystem name match the actual names on the host.

4. Click the **Columns** tab and check the boxes of the columns of data you want displayed in the console. The list consists of several pre-installed columns and of all the columns defined through the Column Definition wizard.

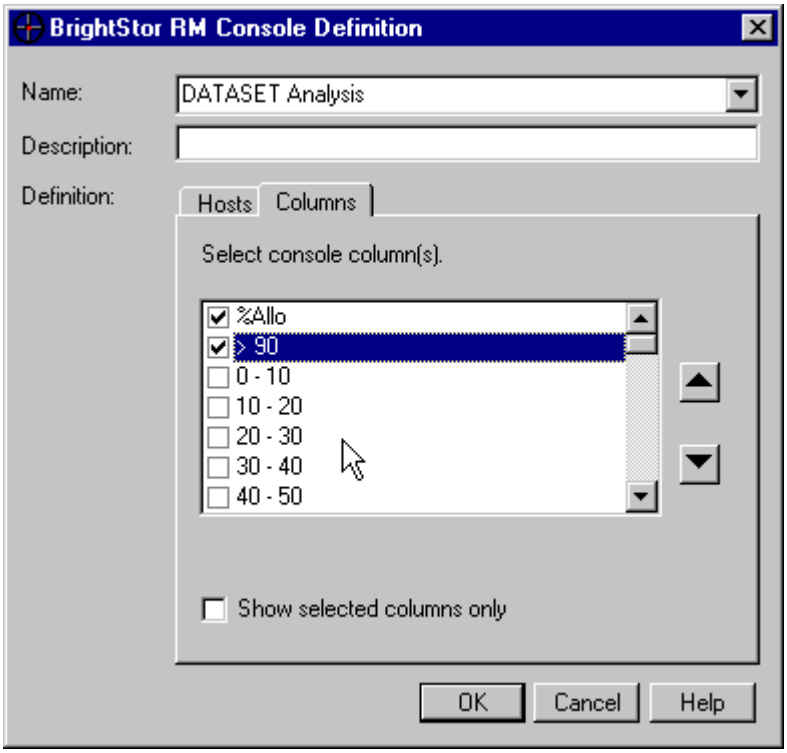

After you make your selections, you can check the **Show selected columns only** box to display only columns of data you have chosen to include in the console.

5. Click **OK**.

The new console is added to the list of consoles that appears in the BrightStor Resource Manager Consoles window.

Selecting Consoles

All BrightStor Resource Manager consoles defined in the system are listed in the BrightStor Resource Manager Consoles window under the **Consoles** icon. To select a console:

✓ Click the **Consoles** icon, then click the desired console.

The selected console is displayed in a console table.

Viewing Consoles

BrightStor Resource Manager consoles are displayed in a console table. The console table shows the hosts and columns of information included in that console.

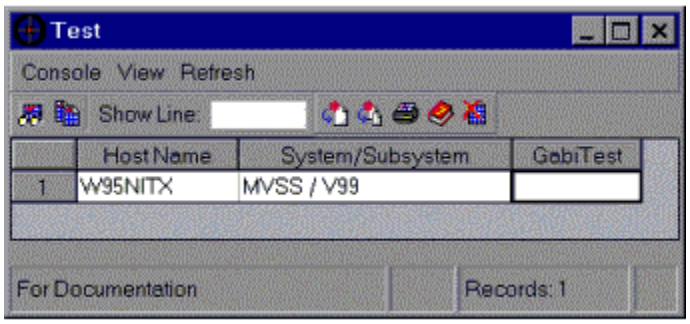

The selected line on the table shows the host that is currently in focus. You can switch between hosts by selecting their corresponding lines in the table. If you open a BrightStor Resource Manager table, it brings data from the host that is currently highlighted on the console.

You can double-click a cell in the console to navigate to its detailed object.

If the console contains columns with global system data, you can right-click on each table cell to obtain the status of the data in that cell. The status display that appears in response indicates whether the cell is waiting for data, the time of the last update, and any errors that may have been encountered during data collection or calculation.

Column Definition Wizard

Use the column definition wizard to define new columns for BrightStor Resource Manager consoles and to modify existing columns. A console column can contain statistics, for example, the number or percentage of volumes that exceed a certain threshold, or it can represent a complex query performed on data collected from several system objects. In all cases, the value that appears in the columns represents a calculated value.

Column Definition Wizard: Consoles

To define a new column:

✓ Click the **Columns** icon in the BrightStor Resource Manager Consoles window and select **New** from the right mouse button menu.

To modify an existing column:

✓ Select the desired column in the BrightStor Resource Manager Consoles window, then select **Definition** from the right mouse button menu.

Both actions invoke the Properties panel of the column definition wizard.

Column Definition Wizard: Properties Panel

Use the Properties panel of the column definition wizard to specify the column name, description, and basic formatting and data collection attributes.

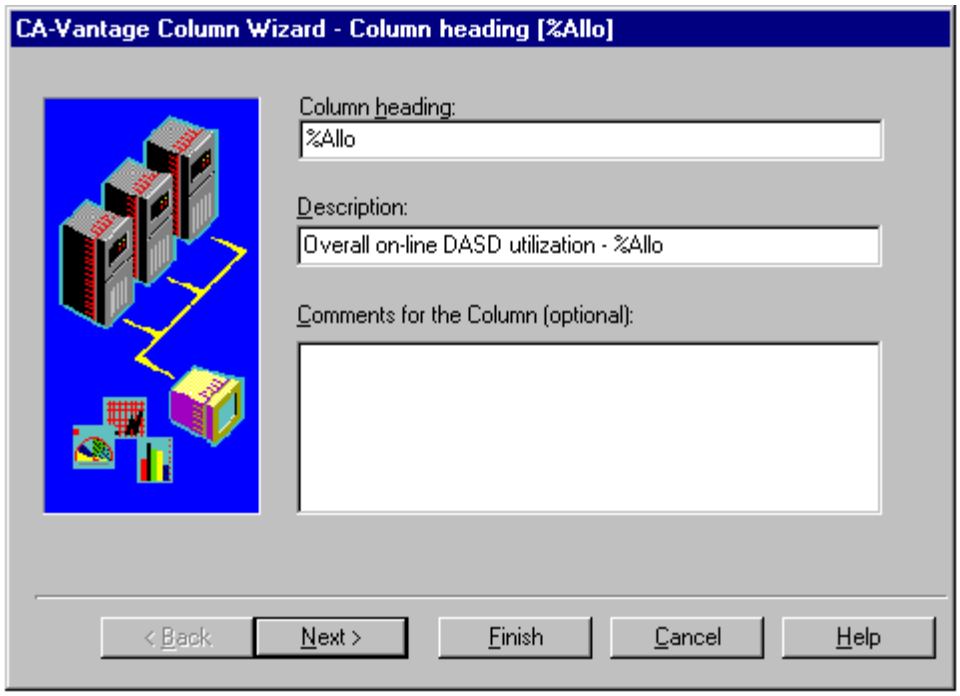

1. Enter the name of the column you want to add in the **Column** heading text box.
2. In the **Description** text box, enter a complete description of the data that appears in the column. The description is shown in the console status line when the cursor is placed on the column heading.

Column Definition Wizard: Object and Filter Definition

Use the Object and Filter Definition panel to select the BrightStor Resource Manager object and, optionally, to apply a filter to it. BrightStor Resource Manager brings the data of the selected object from the host for inclusion in the column.

The value that appears in a console table cell is the calculated result of some function. The value can be based on data collected from several BrightStor Resource Manager objects and derived by means of some algorithm. The value is obtained through the following process:

- A selected range of values, from one or several objects, is brought over from the host and stored locally in one or more Microsoft Access database tables. Use the current panel to select these values.
- A Microsoft Access query is performed on the local database to obtain the value that is placed in the console table cell. The query is defined in Microsoft Access. BrightStor Resource Manager has access to the list of defined queries, from which you can select the one you want to apply to a given column. (A query is provided for each column definition, and a few additional ones are available for general use.) If you want to add new queries to the list of those available, you must define them in Microsoft Access. Use the Data Request panel to select the query.

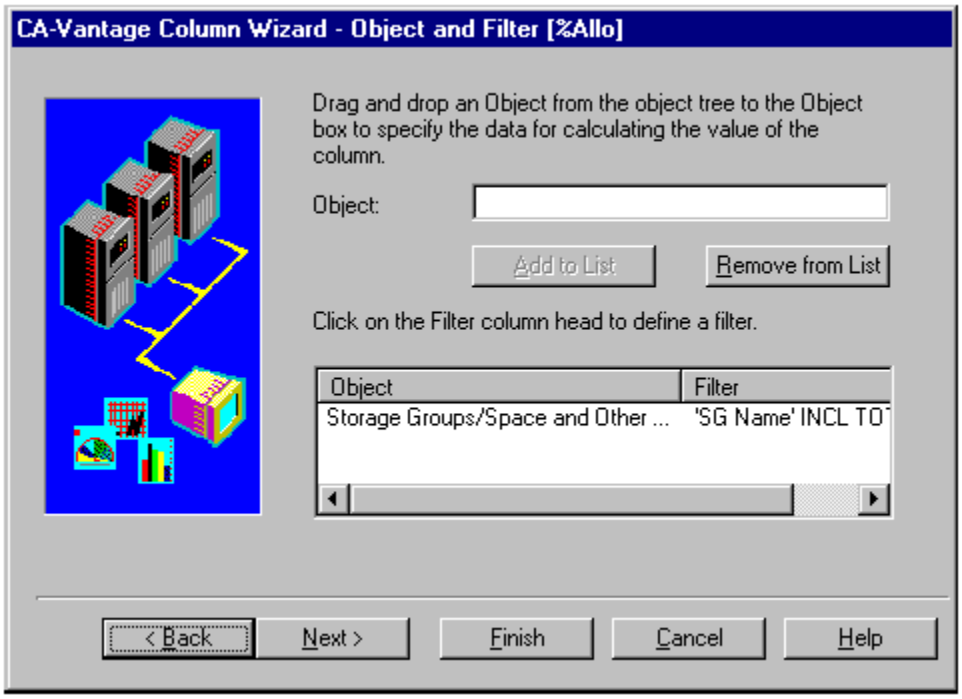

1. From the object tree, drag the object or view you want to use as a source and drop it in the **Source Object** list box. Click the **Add to List** button.

2. To apply a filter to an object, select the object in the list box and click in the column head of the **Filter** column. Use standard filter definition techniques to narrow the range of values to be used from the selected object.
3. Click the **Data** or the **Record Count** button in the **Import Type** section to determine whether actual object data will be collected or only the number of records that meet the selection criteria. If you click **Record Count**, only the number of records in the table is brought from the host, not the actual data contained in the table.

Note: All the rows of the table that pass the filter are selected and used as object data.

You can repeat the operation to add several objects to the list.

The query data extracted from each object is stored in a table in a Microsoft Access database. To view (and change) the name of the database file and of the tables in it:

✓ Check the **Database Information** box.

If the **Database Information** box is checked, clicking **Next** takes you to the Database Information panel. If the box is unchecked, clicking **Next** takes you to the Data Request panel.

Column Definition Wizard: Data Request

Use this panel to select one or more of the predefined queries applied to the Microsoft Access table that stores the object data. The queries extract the desired values from the table and generate the individual value that appears in the console table cell (column).

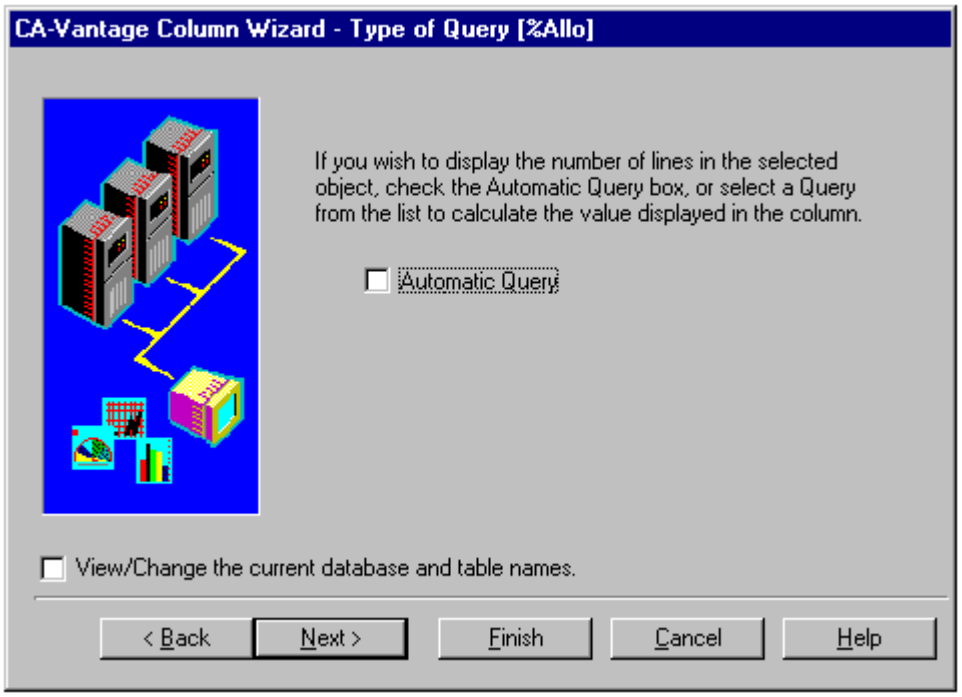

To display only the number of records in the selected object:

✓ Check the **Automatic Query** box. BrightStor Resource Manager displays this number regardless of whether you selected **Data** or **Record Count** as the **Import Type**. No Microsoft Access queries are involved in this case.

Column Definition Wizard: Database Information

The query data extracted from each object is stored in a table in a Microsoft Access database. This panel lists the name of the database file and the names of the database tables that store the various object data. (You need to know the name of this table if you intend to write a new query for this object or to change an existing query). The default database file name is **multi.mdb**.

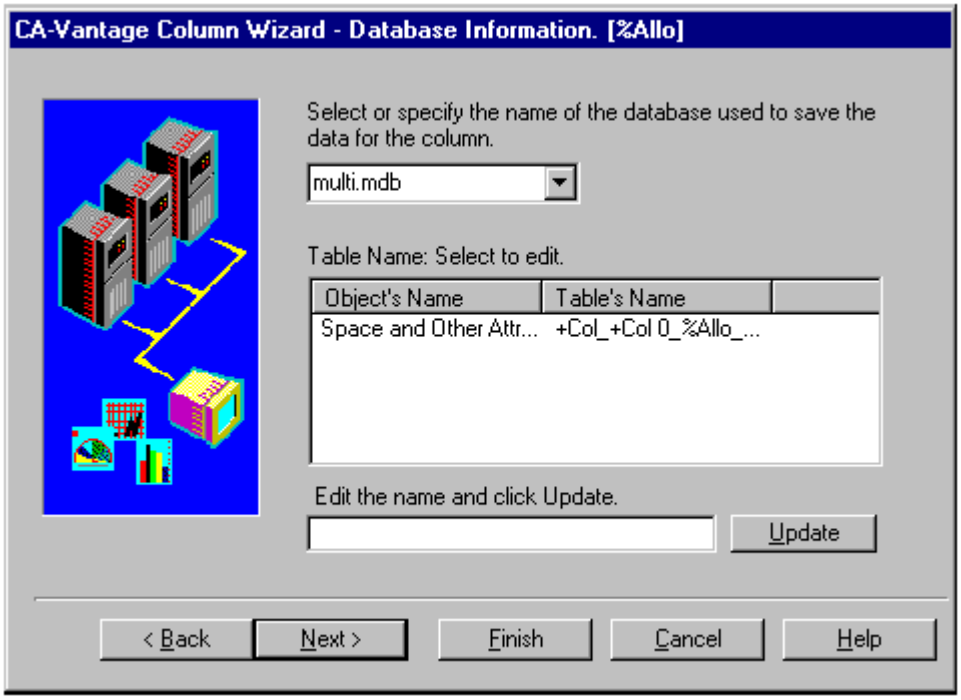

You can change both the database file name and the names of the individual database tables. To change the name of the database file:

✓ Enter the new file name in the **Database** list box.

If the file name you specify does not exist, BrightStor Resource Manager creates it.

To change the name of a database table:

✓ Click on the desired object in the objects list, then enter the new table name in the text box at the bottom of the panel. Click **Update**.

Column Definition Wizard: Color Definition

You can choose to assign colors to ranges of numeric data or to various conditions in the columns. When the panel appears, the default **No Color** button is selected. To define the color of bars in numeric fields or the color in which the entire field is rendered:

1. Click the **Bar Color** or **Full Color** button.

   The display expands into a table with columns for conditions and matching colors. Conditions consist of a logical operator and a value.

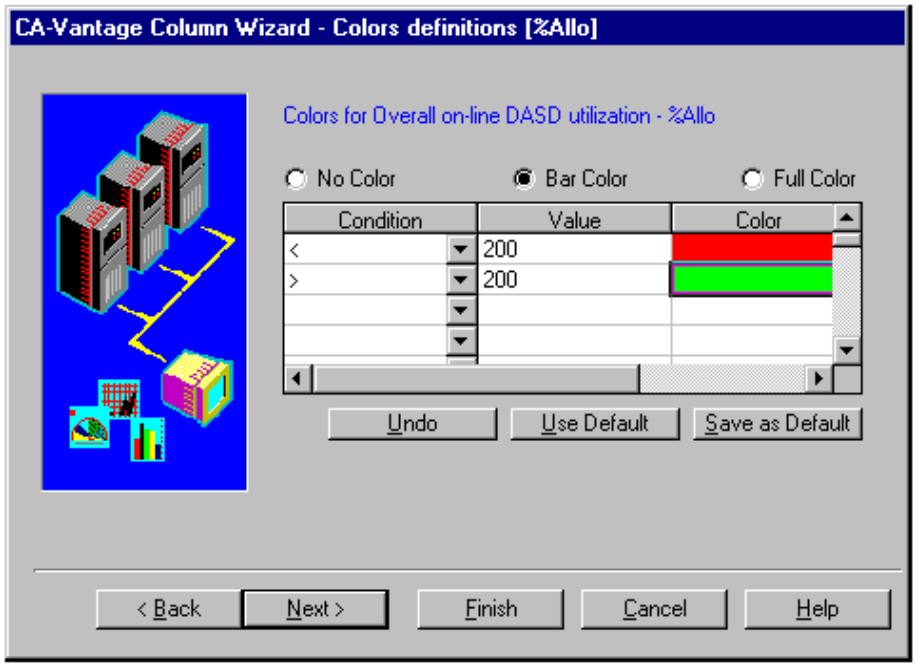

2. Define a condition by selecting a logical operator from the **Condition** column and entering a value in the **Value** column.
3. Assign a color to the condition by clicking the cursor in the **Color** column.
4. When the Color dialog appears, click on the desired color and click **OK**.

5. To define and apply a color other than the ones that appear in the list of Basic colors:

   a. Click the **Define Custom Colors** button.

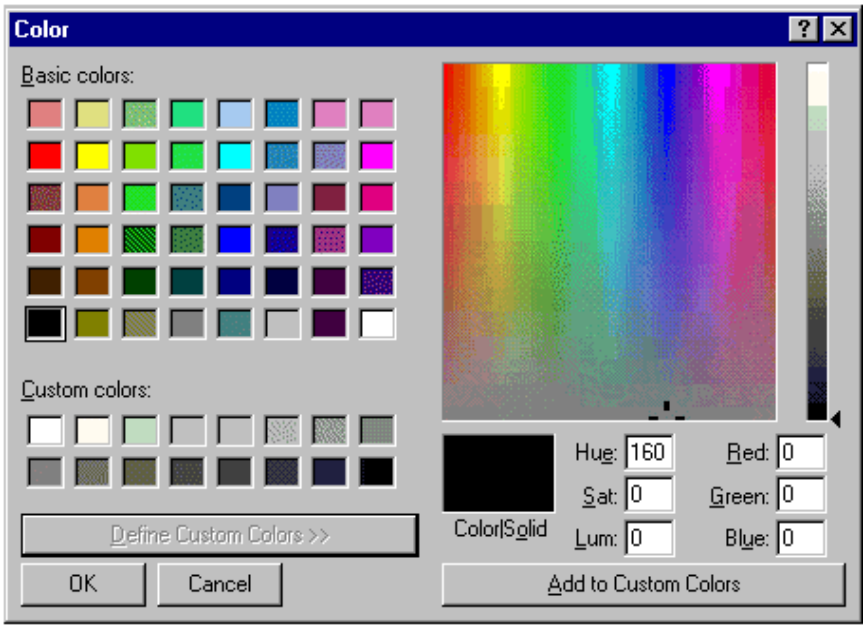

   The dialog expands to the right to let you define custom colors and assign them to one of the 16 squares in the **Custom colors** group.

   b. Select one of the **Custom colors** squares.

   c. Set the **Luminosity** slide bar to the desired level.

      240, at the top of the scale, corresponds to white; 0, at the bottom of the scale, corresponds to black. Note the values changing in the **Lum**(inosity) and **Red**, **Green**, and **Blue** fields as you move the slide bar.

   d. Move the cursor on the color palette over the desired color.

      Note the values changing in the **Hue** and **Sat**(uration) fields. Values in the **Red**, **Green**, and **Blue** fields change within the limits of the selected luminosity.

      Your color selection appears in the **Color|Solid** box.x

   e. Click the **Add to Custom Colors** button when you are satisfied with the selected color.

   The new color appears in the Custom colors group on the left side of the dialog.

6. Click **Next**.

You can apply color definitions you created for one condition to other conditions. To do so:

1. Complete a color definition for any condition.
2. Click Save as Default.
3. Select another condition and click **Use Default**. The definition that has been saved as default is applied to the new condition.

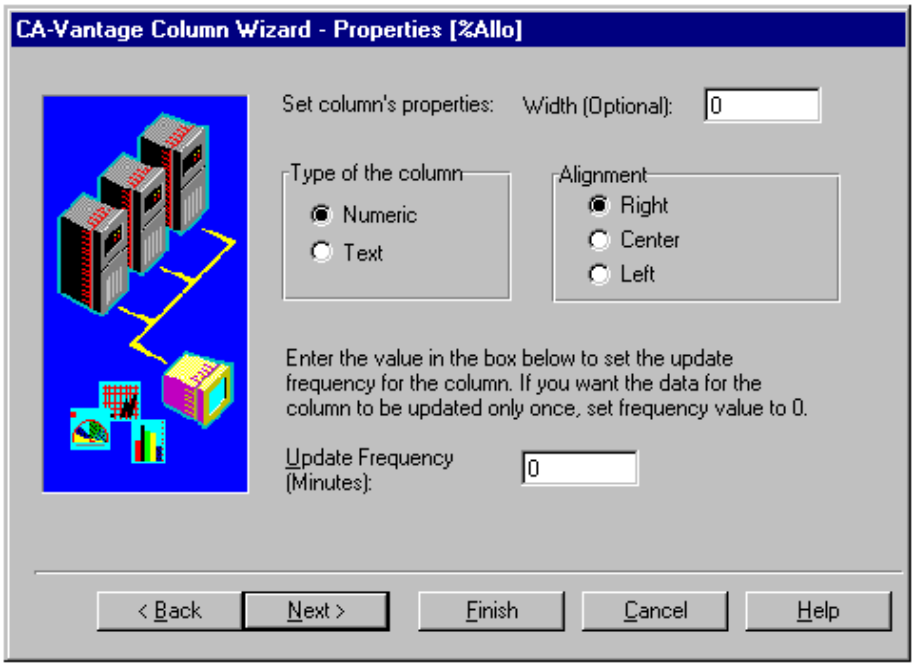

4. Click the appropriate button to specify the alignment of data within the table cell (**Left**, **Center**, or **Right**).
5. Enter a value in minutes in the **Update Frequency** text box to indicate how often the data should be updated in the column that you are defining. If you set this value to 0, data in the column is collected once but not updated.
6. Optionally, the number of bytes used to store the values in this column in the **Width** field. The width of the column that appears in the console table is either the actual width of the heading or the number you provide, whichever is larger.

7. Click **Next**.

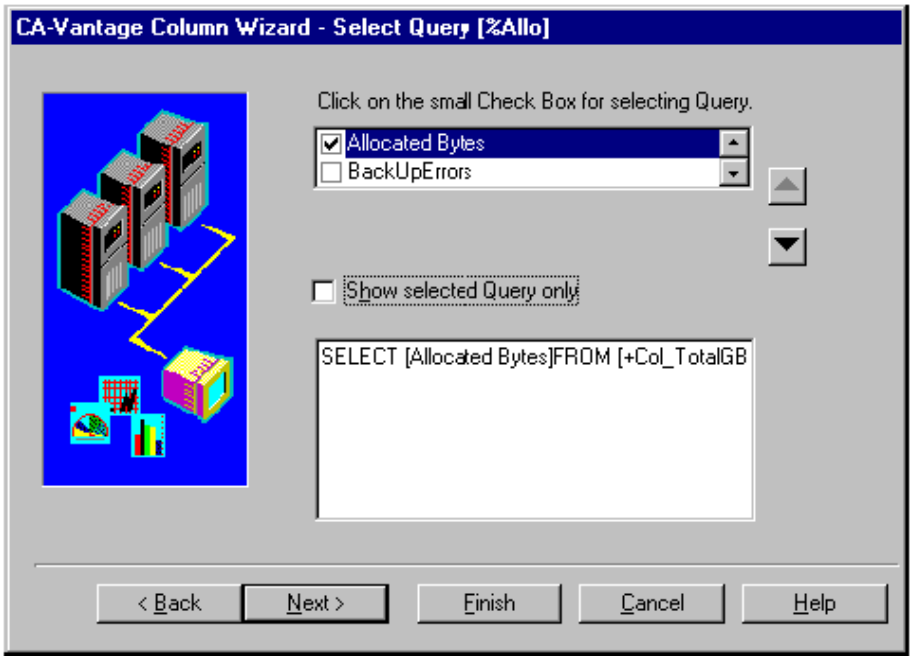

✓ Choose one of the queries listed in the **Query Name** list box. You can select more than one query for a column. When you calculate the value of a cell, all the queries are performed in sequence, and the result of the last query appears in the console table.

After you make your selections, you can check the **Show selected queries only** box to display only queries you have chosen to apply to the object data.

New queries are defined in Microsoft Access. You must be familiar with Microsoft Access to define new queries.

✓ Click **Next**.

Column Definition Wizard: Detail Object

When you double-click a cell on the console display, you navigate to a BrightStor Resource Manager detail object. Use this panel to define the object to which you navigate when you click the cells of the column being defined.

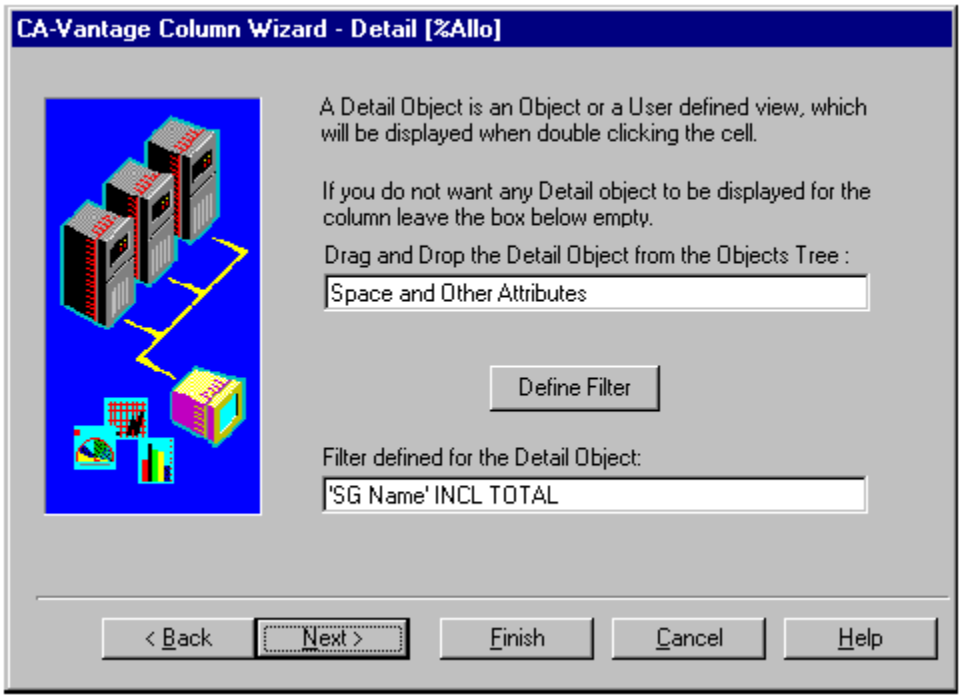

1. Drag the desired object of view from the object tree into the **Detail Object** field.
2. To apply a filter to the selected view, click the **Define Filter** button and use standard filter definition techniques.

Note: If no detail object has been defined for a new column, the last object selected in the Data Request panel for this column becomes the detail object.

3. Click **Next**.

Column Definition Wizard: Alert Definition

Use this panel to define a condition and specify a message to be issued when the condition is met for the given cell.

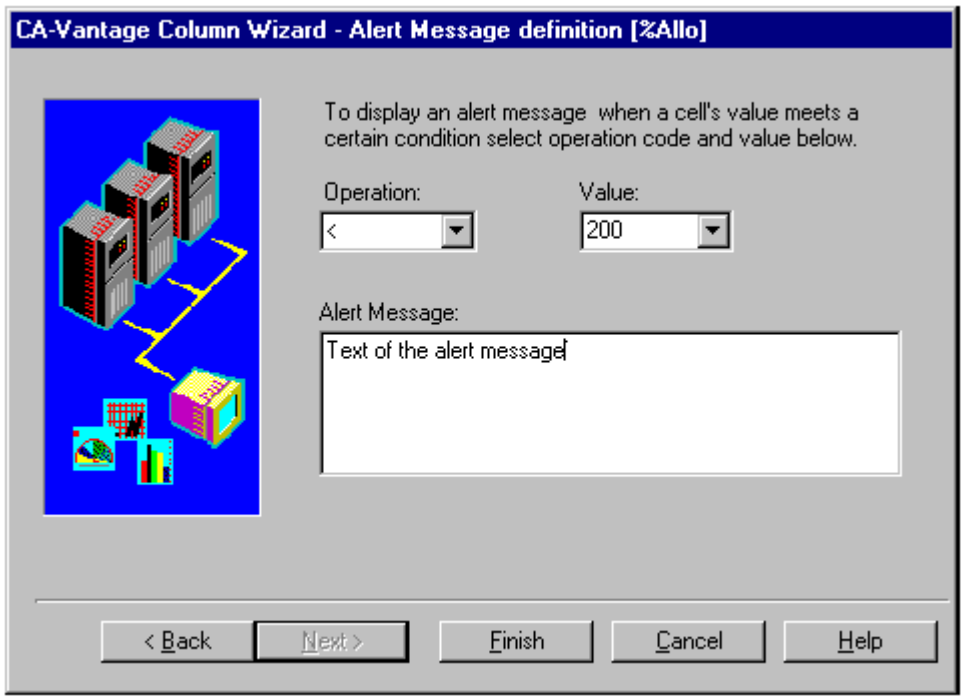

To define an alert:

1. Select one of the **Operation** and **Value** options from their respective drop-down boxes, or enter the desired values manually. For example, if the cell contains free space data in MB, selecting the operator **<** and the value **10** defines the condition **less than 10 MB of free storage space**.
2. In the **Alert Message** text box, enter the message you want the system to issue when the condition is met.

   This is the last panel of the Column Definition wizard.

   To complete the column definition:
3. Click **Finish**.

Index

3

A

B

C

D

E

F

G

H

I

J

L

M

N

O

P

Q

R

S

T

U

V

W

X

Y

Z