Getting Started with
SAS® IT Resource Management 2.7

Getting Started with SAS® IT Resource Management 2.7

Contents

C H A P T E R

1

Introduction to SAS IT Resource Management

Using This Document

Getting Started with SAS IT Resource Management provides an overview of how you can use this product to manage and analyze your IT services. This document supports SAS IT Resource Management software, 2.7 and later.

Chapter 1 — Introduction to SAS IT Resource Management
  This chapter provides general information about this document and concepts about SAS IT Resource Management that will be helpful to new and experienced users. It also presents terms and concepts with which all users should become familiar.

Chapter 2 — Starting SAS IT Resource Management
  This chapter provides information about how to start the SAS IT Resource Management client and server, how to access data on the client and server, and how to use the window interface.

Chapter 3 — Working with Your Performance Data
  This chapter provides information for the SAS IT Resource Management administrator who is responsible for gathering the raw data, creating the PDB, and processing and reducing data into the PDB. In addition, this chapter provides step-by-step instructions for completing the QuickStart, Create PDB, and Process/Reduce Wizards.

Chapter 4 — Analyzing Data through Reports
  This chapter introduces you to various reporting tools that you can use to analyze your data. It also discusses the report styles that are available within these tools.

This chapter is beneficial for all users, but is particularly helpful for analysts who will be creating and running reports on the SAS IT Resource Management client.

Chapter 5 — Additional Documentation and Resources
This chapter provides a list of additional documentation and other resources that you might find helpful while you use SAS IT Resource Management.
    This chapter also provides detailed instructions that enable you to access Web and online information. Web documentation includes the SAS IT Resource Management What's New, User's Guide, Macro Reference, Glossary, and Server Setup Guide. Online documentation includes Collector Updates for this Release, Server Setup Guide, QuickStart Examples, and more.

This document is for SAS IT Resource Management running on UNIX, Microsoft Windows, or z/OS. SAS IT Resource Management provides a graphical user interface (GUI) for UNIX and Windows and a GUI for z/OS. You can also view reports through your Web browser.

Most of the sample windows shown in this document are based on the SAS IT Resource Management GUI for Windows.
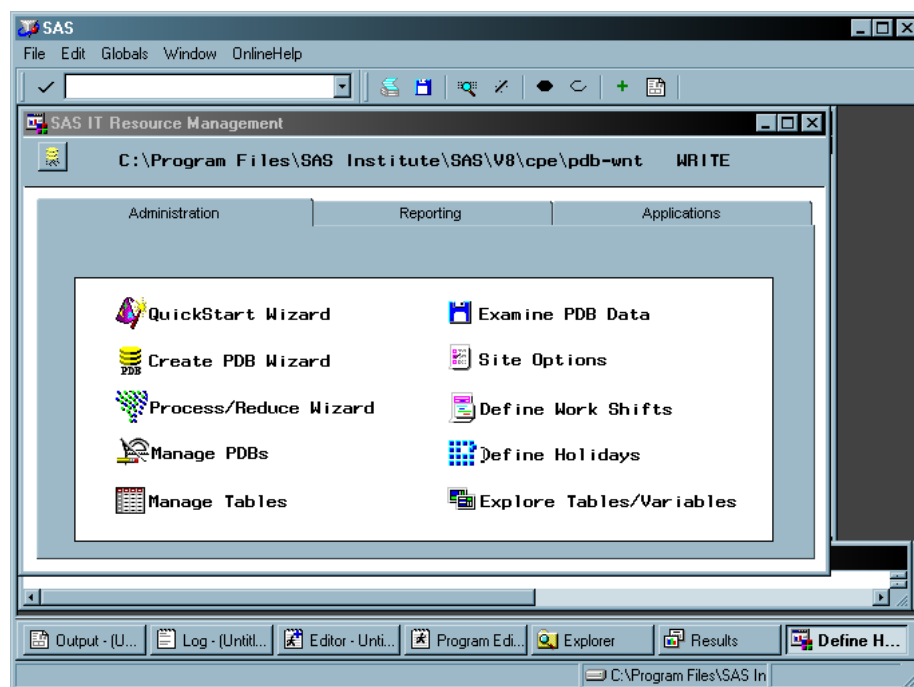
Depending on your operating environment, your windows might differ from the samples shown in this document, especially if you are running on z/OS. Additionally, some of the tasks described in this document must be performed on a specific host, either the client or the server. This is described throughout the document, when applicable.

This document describes how to perform tasks by using the GUIs. However, most tasks that can be performed through the GUIs can also be performed in batch mode by using SAS IT Resource Management batch macros. For example, you can start the application in batch, process and reduce data in batch, and create and run reports in batch.

The purpose and syntax for each macro are available in the SAS IT Resource Management macro reference documentation. For information about how to access

documentation about the macros, see Chapter 5, "Additional Documentation and Resources," on page 65.

SAS IT Resource Management Overview

SAS IT Resource Management is a performance evaluation tool that helps you to utilize your IT resources more effectively. It enables you to access, manage, and analyze large quantities of performance data. Performance evaluation is expanding to include areas such as troubleshooting, planning for growth, Web servers, applications support, and phone systems. SAS IT Resource Management provides the technology that lets you summarize large volumes of data into a common format and then quickly and easily analyze the data by using the reporting tools within SAS IT Resource Management.

By using SAS IT Resource Management, you can read and process data from almost any data source. SAS IT Resource Management supports many popular network and systems management tools and data sources on z/OS, UNIX, and Windows. It also enables you to read and analyze data from any time-based data source that contains a datetime stamp and whose log record formats are known.
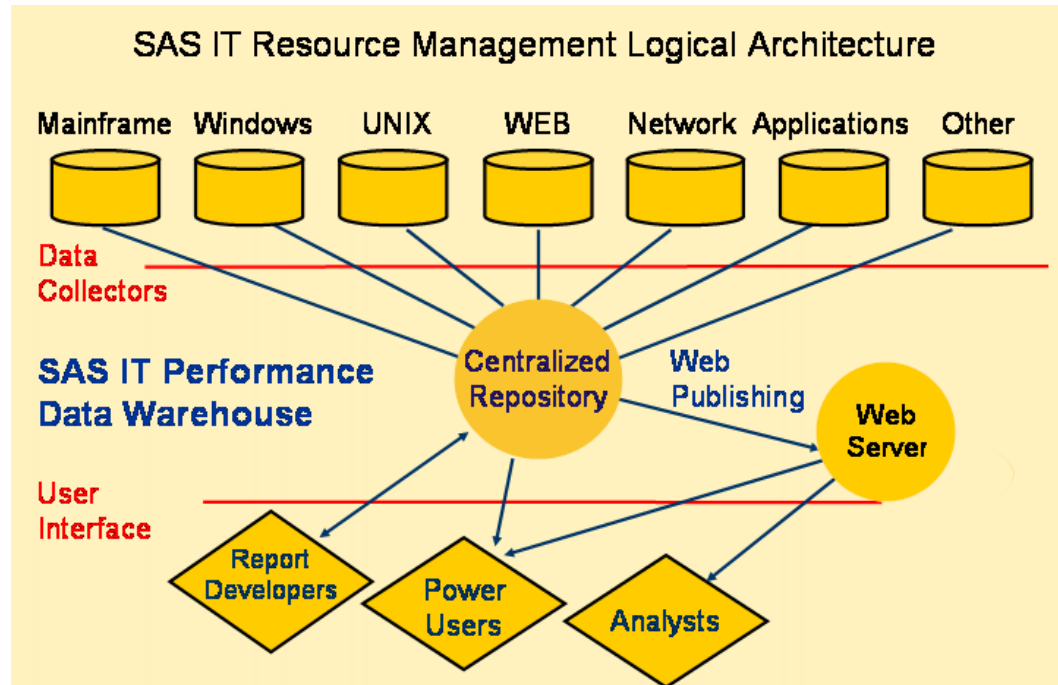
The software that you use to gather data about your IT services is referred to as a *collector*. SAS IT Resource Management supports many popular data collection applications by providing tables for reading and storing the data. However, you can also create your own custom collector, and SAS IT Resource Management can read and process that collector's data into your custom tables.

SAS IT Resource Management reads raw data (metrics) in the format that they are logged by your collector. SAS IT Resource Management can minimize the volume of data that you keep, by summarizing the detailed performance data into smaller groups or levels in the performance data warehouse (PDB). You select the statistics to be calculated at each summary level. After the data is reduced into the summary levels in

the PDB and the statistics are calculated, you can analyze information over different time periods by running reports on a specific summarized group of your data.

You can also customize many aspects of the SAS IT Resource Management application. You can use or modify tables and reports that are supplied with SAS IT Resource Management and you can create new tables and reports.

How Can SAS IT Resource Management Help You Manage Your IT Resources?

SAS IT Resource Management can be used in a variety of systems management situations within your organization. You can use SAS IT Resource Management as follows:

- □ to establish baselines for system use and performance. You can then track performance and workload volume against these values to anticipate the need for resource upgrades or to identify the causes of unsatisfactory service.
- □ to track activity and to highlight trends in usage and performance.
- □ to analyze your IT systems to determine the optimum resources required, so that you spend only the money and resources that are necessary to maintain excellent customer support.
- □ to run exception reports that can identify problems in your IT systems.

SAS IT Resource Management features can help you to make effective business decisions concerning your IT services. A few of these features are as follows:

- □ You can process data from almost any data source.
- □ From a client, you can access and analyze data stored on a local or remote server.
- □ You can view reports through a SAS IT Resource Management GUI or through your Web browser.
- □ You can summarize data in the PDB, thus reducing your storage needs. You can then analyze and create reports based on the summarized data.
- □ You can point and click to access and navigate around your PDB and reports.
- □ You can work with demonstration PDBs for several popular data collection applications.
- □ You can create custom or drill-down reports with a variety of reporting tools.
- □ You can modify and create rules that identify exceptions in the data.
- □ You can access examples that set up various data sources for use with SAS IT Resource Management.
- □ You can perform tasks in batch mode by using an extensive library of SAS macros. You can also perform tasks interactively through a SAS IT Resource Management GUI.
- □ You can access the extensive analysis capability and support provided in the SAS System.

What Is a Performance Data Warehouse?

A *performance data warehouse* is a group of SAS libraries that contains the data used by SAS IT Resource Management. (A performance data warehouse can also be called a performance database or PDB.)

The PDB contains data libraries, tables, and variables that define which statistics are calculated for the data. The PDB also stores options associated with the data.

Your site can have one or more PDBs, depending on how you want to collect and store your data. These are a few of the reasons why you might decide to employ multiple PDBs:

☐ You can use a separate PDB for each data source, such as Windows servers, UNIX servers, and phone switches.

☐ Different groups might want to manage their own separate PDBs.

☐ You might want to create a test PDB so that you can investigate specific performance problems or facilitate testing.

Alternatively, you can choose to combine data from multiple data sources in the same PDB. You can organize your data and PDBs to meet your site's needs.

SAS IT Resource Management supplies many table and variable definitions. These supplied definitions are the default specifications that govern how these tables or variables are created. When you create a PDB at your site, you can perform the following tasks:

☐ Add the supplied table and variable definitions and use them in your new PDB.

☐ Add the supplied table and variable definitions to your new PDB and modify them for your specific needs.

☐ Create new table and variable definitions in your PDB.

For example, if you add a supplied table and its supplied variables to your PDB, you might want to change defaults such as the age limits of data in each data level, the settings that specify which data is kept at detail level, and which statistics are calculated at each summary level. You might also want to delete variables that you do not want to use or to create new variables in order to calculate specific statistics.

The task of setting up your PDB is described further in "Using the QuickStart Wizard" on page 21.

Tables, Variables, and Definitions in a PDB

A *table* is a group of data, statistics, and definitions that is stored in a PDB. Each table contains data related to a specific performance area, such as a specific type of network activity. Each table also has a *table definition*, which is a description of the table's properties.

Table definitions include:

☐ The name of the collector that is used to gather the data

☐ The type of table — interval or event

☐ The age limits of data in each PDB level

☐ The list of variables that are stored in the table.

Variable definitions are also associated with a table. A *variable definition* is a description of a variable's properties. Variable definitions include:

☐ The data type, which defines whether the variable is numeric or character

☐ The interpretation type, which is used to determine the default statistics kept for the variable and how the variable is calculated in the summary levels

☐ The format, for external display

☐ The length, for internal storage

☐ The statistics to be calculated for the variable at each summary level in the PDB.

You can use a GUI or batch mode to add supplied table definitions to your PDB. You can also modify table definitions and create new ones interactively or in batch. The

table and the variables within the table are created automatically by SAS IT Resource Management from the definitions.

You can also create definitions for formula variables and derived variables. *Formula variables* are variables whose values are calculated from the data that you collect. These variables are calculated dynamically when you access the data so that you do not need to store the values of these variables in the PDB. *Derived variables* are also variables whose values are calculated from the data that you collect. After they are computed, however, the values of the derived variables are stored with your data in the PDB. SAS IT Resource Management does not have to calculate them again each time they are requested.

Within a table, data is stored in five physical libraries (known logically as levels within a table): DETAIL, DAY, WEEK, MONTH, and YEAR. The data in the detail level is typically very similar to the data as it was recorded by your collector. The other four levels, also known as reduction or summary levels, contain your reduced or summarized data. The day level contains data summarized by the day, the week level contains data summarized by the week, and so on. For example, a table for physical disk data can contain detailed information for each physical disk in the detail level and also summarized data in each of the four summary levels (day, week, month, year).

When you display the data from one level of a table in the PDB, this "picture" of the data is referred to as a *view*. A view is named by using a combination of the level and the table name in the form of *level.table*. For example, the view for the detail level in a table named NTPHDSK would be DETAIL.NTPHDSK. The view for the day level of the same table would be DAY.NTPHDSK. When you report on your data, the reports are based on the views.

On UNIX and Windows, you can see a list of tables in a specific PDB, as well as status information about the tables, by selecting `Manage Tables` from the Administration tab in the main window of SAS IT Resource Management. To see a list of variables for one of the tables, select a table from the list and select this menu path: **Locals ▶ List Variables**.

On z/OS, you can see the corresponding information by selecting **PDB Admin ▶ Config Active PDB Dictionary ▶** *select a table* **▶ ItemActions ▶ List Variables**.

Summarizing Data into Libraries in the PDB

As previously mentioned, data is stored physically in five data libraries within tables. These libraries are known logically as *levels*. When your data is first transferred into the PDB from your raw data file, this is known as processing your data into the PDB, or the *process* task. This processed data is stored in the detail level. The data in the detail level is very similar to the original data that you collected, with simple transformations (when applicable), such as converting continuously ascending counters into rates per second or adjusting a counter that has exceeded its maximum and restarted at zero. For collectors with vast amounts of data, such as transaction records, the process task can also perform some pre-summarization.

After the data has been processed into the detail level of the PDB, you then summarize or reduce the data into the other levels of the PDB. These other four levels are called *summary levels* or *reduction levels*, and they are day, week, month, and year.

Logical Data Flow
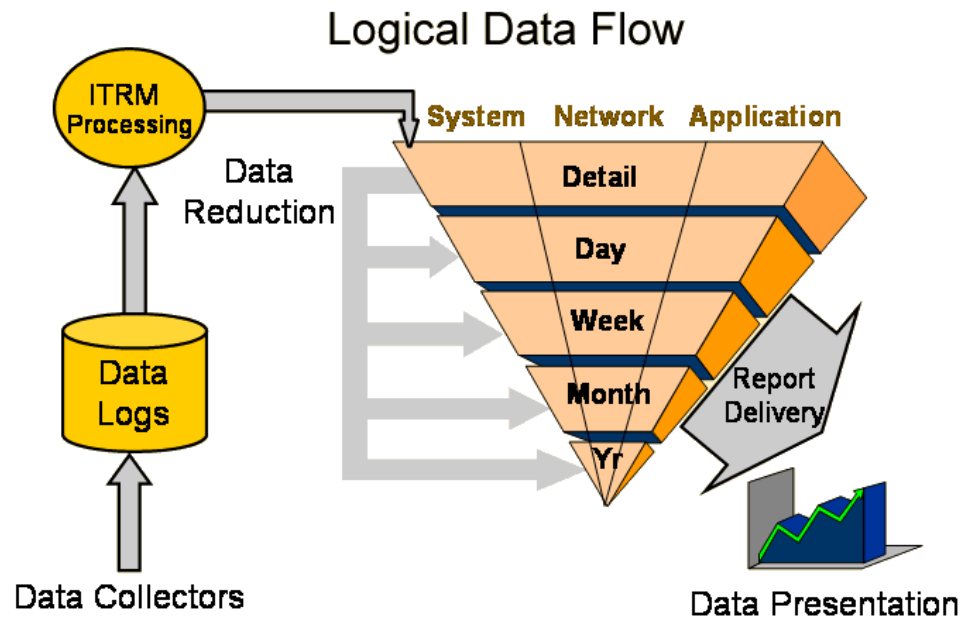
The DETAIL library
> contains your data after you process it into the PDB. Each observation in a data set represents an event or an interval. The data in the detail level is very similar to your data records in the raw data file.

The DAY library
> contains daily statistics summarized by grouping variables, such as machine, hour, and day. These statistics are calculated from the table's data in the detail level. For example, each observation in a data set in this library can contain statistics for a specific machine, for a specific hour of the day, for a specific shift, and for a specific day.

The WEEK library
> contains weekly statistics summarized by grouping variables. The statistics are calculated from the table's data in the detail level. For example, each observation in a data set in this library can contain the statistics for a specific machine, for a specific hour of the day, for a specific shift, and for a specific week.

The MONTH library
> contains monthly statistics summarized by grouping variables. The statistics are calculated from the table's data in the detail level. For example, each observation in a data set in this library can contain the statistics for a specific machine, for a specific hour of the day, for a specific shift, and for a specific month.

The YEAR library
> contains yearly statistics summarized by grouping variables. The statistics are calculated from the table's data in the detail level. For example, each observation in a data set in this library can contain the statistics for a specific machine, for a specific hour of the day, for a specific shift, and for a specific year.

When data is "reduced" or summarized, it is copied from the detail level and reduced into the summary levels in the PDB. (These summary levels are independent of each

other.) The data is reduced based on the statistics that you select for each summary level. You can use the default statistics, or you can customize the list of statistics that you want to be calculated for your data at each summary level. When the data is reduced, the specified statistics are calculated, and these summary statistics (not each observation from your data) are stored at each summary level.

Note: You do not have to use all levels of a table. Furthermore, depending on your requirements, you can have different variables and statistics at each level. △

Note: Be sure to make a backup copy of your PDB before you make changes to it, and again after you have customized it. △

How Long Is Data Stored in the PDB?

The detail library might easily become very large because this library contains the data as it is logged by your performance management tool. As you process new data into the PDB and reduce the data into the summary levels, you will want to decide how long to keep data at each level in the PDB.

To reduce the space required for the detail library and the other libraries, your site can set age limits on the data at each level. After data is reduced from the detail level into the summary levels, you might not need to keep the detail level data, or you might want to keep only a few days of data at that level. Similarly, for the other summary levels, you can choose a time limit, such as 3 days, 10 days, a month, and so on. The data will be deleted from each level when its age limit is reached. By using each of the summary levels, you can still analyze and report on long-term trends for your IT systems, but you do so by using summarized data instead of large quantities of detail data.

The following built-in age limits are assigned by default to the levels in the table:

□ Detail — 10 days
□ Day — 45 days
□ Week — 15 weeks
□ Month — 18 months
□ Year — 5 years

You can change the age limits by following this path from the `Administration` tab within the GUI for UNIX and Windows: **Manage Tables** ▶ *select a table* ▶ **File** ▶ **Properties**. From the Table Properties window, select the `Table Status` tab to review the age limit settings.

The corresponding path in the GUI for z/OS is **PDB Admin** ▶ **Config Active PDB Dictionary** ▶ *select a table* ▶ **ItemActions** ▶ **Edit Definition**.

SAS IT Resource Management Libraries

SAS IT Resource Management uses several libraries that contain items such as software, data, user options, site options, and more. These libraries are described below.

The Program Library

The SAS IT Resource Management program library (PGMLIB) is a SAS library that contains the SAS IT Resource Management software for your current host. It contains

the master data dictionary (supplied table and variable definitions), supplied report definitions, supplied rule definitions, and the SAS IT Resource Management client or server software programs for your current host. Supplied table, variable, report, and rule definitions are provided for many popular data sources to assist you in using those collectors and data sources with SAS IT Resource Management.

When you create a new PDB, it contains many default settings. Your site can customize those default options, customize the supplied definitions, create new definitions, and add table and variable definitions to your PDB from the master data dictionary (in PGMLIB). You can also install user-generated table definitions in the master data dictionary, and your site can easily add those definitions to other PDBs at your site.

When SAS IT Resource Management is updated in a new release, the information in PGMLIB is updated before shipping. You cannot modify PGMLIB. However, you can access PGMLIB to view, search, or copy from the list of supplied table definitions, variable definitions, report definitions, and rule definitions.

The CPSYSLIB Library

The CPSYSLIB Library is a SAS library that is used on UNIX and Windows. The library contains the programs, screen definitions, and client-related information for the SAS IT Resource Management GUI for UNIX and Windows.

It is not available on z/OS.

The Site Library

The Site Library contains the site-related information, such as site holidays, default work shift schedules, default start of week setting, and more. When you or others at your site create a new PDB, the default properties for that PDB are copied from the site library. After you create a new PDB, you can customize the settings for that PDB.

A default site library is created when you install SAS IT Resource Management. Therefore, a site library is installed on your client hosts and a site library is installed on your server host. If you are processing data for multiple sites or groups within your organization, then you might have a separate SITELIB library for each site or group. Thus, you can customize the default settings for each group.

In the GUI for UNIX and Windows, you can view many of your site options from the `Administration` tab in the main SAS IT Resource Management window. To view the name of the currently defined site library, select `Site Options` from the `Administration` tab. Select `Change Sitelib` to display the name of the current site library. You can modify the current site library from this window.

The corresponding path in the GUI for z/OS is **PDB Admin ▶ Define Work Shifts ▶ Site Defaults**.

If your organization wants to change the defaults in the site library to customize one or more options, you should copy the default site library and make changes to the copy. To modify settings in the site library, you must have write access permission for this library. However, you can read and copy items from the active site library if you have read access to the library. For more information about accessing your site library or setting site-specific preferences such as holidays and shifts, refer to the *Work with a SAS IT Resource Management Site Library* chapter in the SAS IT Resource Management User's Guide. For information about accessing the User's Guide, see Chapter 5, "Additional Documentation and Resources," on page 65.

Note: Be sure to make a backup copy of your SITELIB before you make changes to it, and again after you have customized it. △

The SASUSER Library

The SASUSER library is a SAS data library that is used by your interactive SAS session for storing user-specific information. Examples of this information are the locations of PDBs that you use and any customizations that you specify for SAS options.

When you start the SAS System, a SASUSER library is automatically created for you, if one does not already exist. When you customize keys, menus, or other items within the application, these modifications are automatically saved in your SASUSER library. By default, any source, log, or report files that you save are saved in this library.

Note: Be sure to make a backup copy of your SASUSER library before you make changes to it, and again after you have customized it. △

By default, this library is not used in batch mode on z/OS. For platform-specific information about the SASUSER library, refer to the SAS Companion documentation for your platform and your current version of SAS.

The Performance Data Warehouse (PDB)

In addition to the data libraries described in "What Is a Performance Data Warehouse?" on page 4, the PDB contains several other libraries related to the data in a specific PDB. These libraries must be present in a PDB, even if they are empty.

The ADMIN library
is available to all PDB users. Therefore, it might be a convenient place to store graph and text reports, report palettes, logs, and so on, within a PDB. If you are using the data filtering routines to prevent duplicate data when you process your data into the PDB, this library is used by those routines.

The COLLECT library
is used for temporary data storage while your data is being processed into the PDB. For example, you can use the COLLECT library to store any staging data sets that you create for use by the Generic Collector Facility software. This library can also contain information or libraries used with MXG software.

The PDBWORK library
is used by the reduction step for intermediate data. The PDBWORK library (which is permanent) is used instead of the WORK library (which is temporary) so that reduction can be restarted from a checkpoint if necessary. Sometimes PDBWORK contains temporary SAS data sets, which are used for calculating statistics for the day, week, month, and year levels.

The data dictionary
is in each PDB. Also known as the DICTLIB library, the PDB's data dictionary contains general PDB options as well as information about the data that you have stored in each level of the PDB. For example, the DICTLIB library contains specifications about which data to keep in the PDB, and a count of the number of observations in each level of the PDB.

Additional Libraries

Depending on your host, additional items can be stored in other libraries.

☐ On UNIX and Windows hosts, additional software can be stored in an auxiliary library.

□ Items such as MXG formats can be stored in an auxiliary library.

□ If you have more than one release level of SAS IT Resource Management installed, you will have a separate version of the PGMLIB library for each release.

□ If you elect to archive data, archive libraries are created.

□ When you run the QuickStart Wizard, the Wizard creates additional storage locations for materials, such as programs.

The Data Collection and Analysis Process
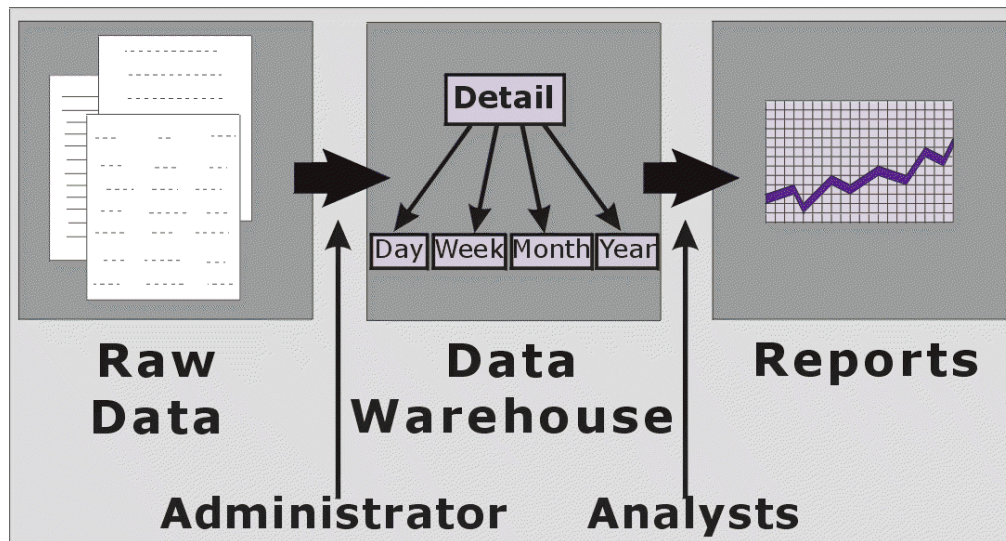
The routine tasks that you will perform to process and reduce data and view reports can be summarized as follows:

1 *Set up your data collector to work with SAS IT Resource Management.* You can use the QuickStart Wizard to create your PDB, configure your site's PDB options within SAS IT Resource Management, and create batch jobs that can be used to process data, reduce data, and create reports based on your data.

   If you use the QuickStart Wizard, the process, reduce, and even the reporting tasks mentioned in the following steps can be performed through daily jobs that are created when you run the QuickStart Wizard. Based on the data category you select in the Wizard, the Wizard also provides an instruction file to assist you in setting up your collector with SAS IT Resource Management. For information about using the QuickStart Wizard, see "Using the QuickStart Wizard" on page 21.

   Whether or not you use the QuickStart Wizard, information about setting up your data collector is available in "Part 2: Setup" of the SAS IT Resource Management User's Guide.

2 *Collect the IT data that you want to analyze.* You can collect data by using an existing performance management tool, or by using a custom tool that you write for a specific data source.

3 *Process raw data* into the detail level of the tables in your PDB. Your site can create a daily job to read and process the data, or you can use the daily job that is created when you run the QuickStart Wizard.

4 *Reduce the detail-level data* into one or more summary levels in the PDB: day, week, month, or year. In most cases, you or an administrator at your site will use

a job generated by the QuickStart Wizard, or will set up a daily job to reduce the data immediately after it is processed into the PDB.

5 *Run supplied report and rule definitions or create and run your own custom report and rule definitions* to help you analyze your data and identify performance issues. You can generate reports by using data from any level of the PDB.

2
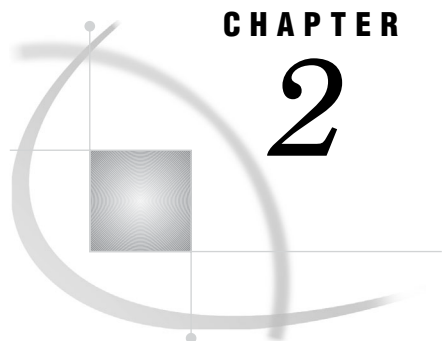
Starting SAS IT Resource Management

Accessing Data on the Client and the Server

SAS IT Resource Management software provides a client/server interface that enables you to process and summarize your data into the PDB on a server, and then analyze and create reports on the client.
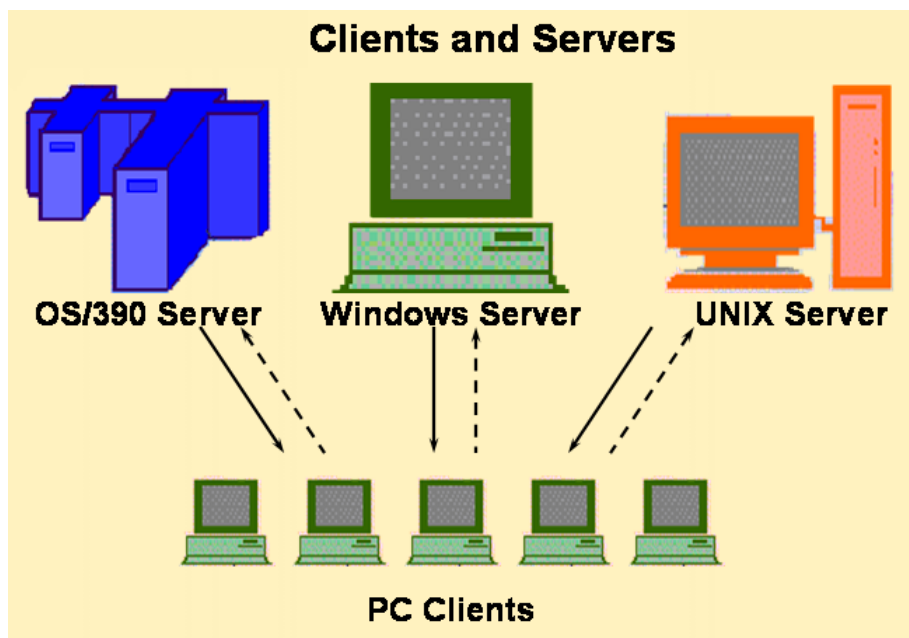
From the SAS IT Resource Management server, you can access local PDBs in both read and write mode. This means that when you are working on a host where SAS IT Resource Management server software is installed, and you have write access to a PDB and all of the associated files, libraries, and subdirectories, then you can open, read, modify, and save changes to the PDB. For example, you can create or modify a PDB, set site options, and/or process and reduce your IT service data. Within the user interface,

these types of tasks are accessible from the `Administration` tab (on UNIX and Windows) and the `PDB Admin` and `PDB Data` buttons (on z/OS).

From the SAS IT Resource Management client, you can access, read, and run report definitions by using data in a PDB that is stored on the server. From the client, you have read-only access to the PDB and limited write access to other related items. For example, from the client, if you have write access to the location where you want to save the items, you can save report definitions, report folders, palettes, reports, rule definitions, and rule folders on a remote server.

Note: To access a remote PDB from the client, you must create and assign a Remote Server Profile, which identifies things such as how to access the PDB and where the PDB is located (see "Analyzing Data on a Remote Server" on page 41). △

Starting SAS IT Resource Management

You can start SAS IT Resource Management several ways. The method that you use might vary depending on your site.

`On Windows:`

To start the SAS IT Resource Management GUI on Windows, follow this menu path: **Start ▶ Programs ▶ SAS IT Resource Management**.

`On UNIX or Windows:`

To start the SAS IT Resource Management GUI on UNIX or Windows, follow these steps:

1 Start the SAS System.

2 On the SAS command line or in the command bar, type `itrm` or `itsv`.
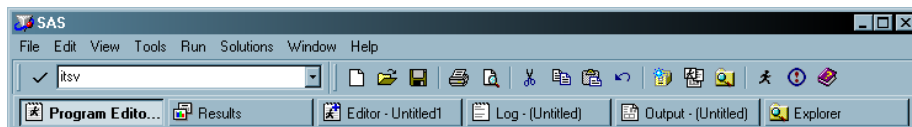
3 Press `Enter` to submit the command.

Note: You can use the above steps for both client and server software on UNIX or Windows. △

`On z/OS:`

To start the SAS IT Resource Management server on z/OS, follow these steps:

1 Start the SAS System.

2 Start SAS IT Resource Management from within SAS. The method used to start the product might vary depending on your site:

□ You can submit a call to the %CPSTART macro (%CPSTART;) from the SAS Program Editor window or in batch mode.

□ You can also type *CMCPE* from the TSO READY prompt to use the clist *CMCPE* that is in the SAS IT Resource Management PDS named *CPMISC*.

The steps for starting SAS IT Resource Management might vary slightly depending on your operating system and on any customized scripts that your site might be using to start the SAS System. For more information, refer to the SAS companion for your current host.

Note: On *UNIX or Windows*, the first time that you start SAS IT Resource Management, you are prompted either to use a demonstration PDB that is shipped with this product or to create a new PDB.

To use a demonstration PDB on UNIX or Windows, select the **Open Existing PDB** option, select a PDB from the list and select **Activate**. Select **Close** to close the Manage PDBs window and return to the main SAS IT Resource Management window.

To create a new PDB, you can use the QuickStart Wizard or the Create PDB Wizard. For more information about creating a PDB, see Chapter 3, "Working with Your Performance Data," on page 19.

On *z/OS*, the first time that you start SAS IT Resource Management, you are prompted for the name of a PDB, only if you have not specified one on the PDB= parameter of the %CPSTART macro. Typically, z/OS customers start SAS IT Resource Management with the PDB= parameter already coded. If you do not, you can specify the PDB named *high.level.qualifiers.*PDBDEMO, where *high.level.qualifiers* are the ones used for this PDB at your site.

A record of all the PDBs that you access is stored in your SASUSER directory. Therefore, after your first use of SAS IT Resource Management, the previously activated PDB is activated, if you start the software without specifying the PDB= parameter. △

Overview of the Window Interface

The SAS IT Resource Management GUI for UNIX and Windows provides a windows environment with easy access to your data and reports. The main window displays tabs (file folder-style) that provide easy access to functionality within the product. The **Administration**, **Reporting**, and **Applications** tabs are shown in the SAS IT Resource Management window below.
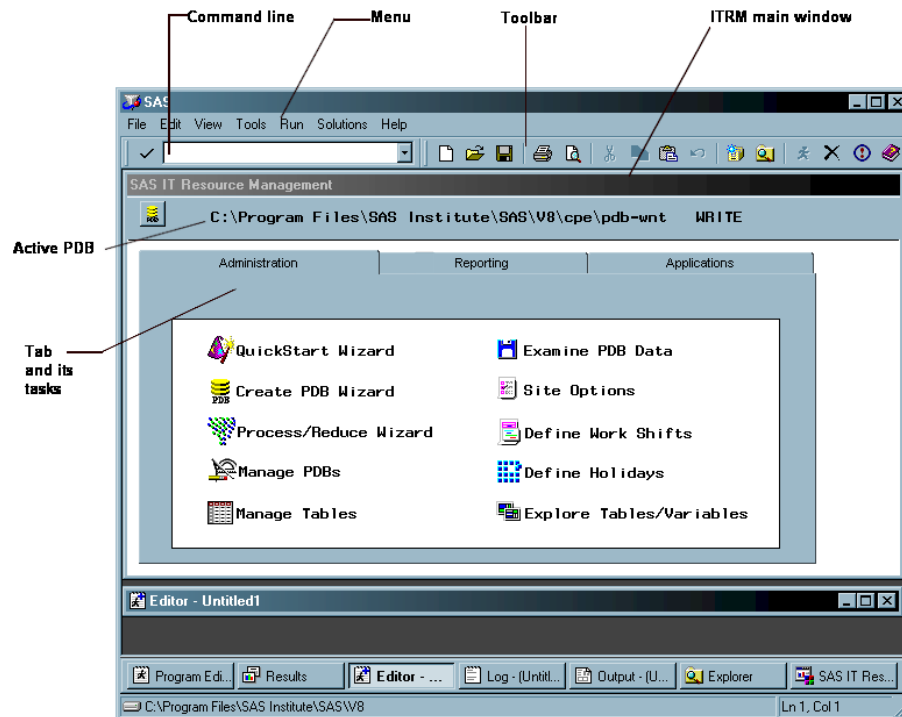
Administration
> contains tasks related to processing and reducing data, and creating and managing your PDB, tables, variables, and site library.

Reporting
> lists several unique reporting tools that you can use to analyze your data.

Applications
> enables you to store and access user-, PDB-, and site-specific applications.

To display the items from one of the tabs, select the tab in the main SAS IT Resource Management window.

The SAS IT Resource Management GUI for z/OS has three rows of buttons corresponding to the three tabs.

The main window also displays the name of the active PDB, your access level for that PDB, and the toolbar, which provides easy access to common tasks for the current window.

To activate a new PDB, follow these steps in the GUI for UNIX and Windows:

1 Select **Manage PDBs** on the **Administration** tab.

2 From the list in the Manage PDBs window, select the PDB that you want to activate.

3 Select **Activate**.

4 From the Activating PDB window, select your access level (read or write) and indicate whether the PDB is local or remote.

5 Select **OK** to close the window and then select **Close** to return to the main window.

In the GUI for z/OS, the corresponding path is

1 Press **PDB Admin**

2 Select **Manage PDBs/Switch PDB**

3 Select the PDB that you want to activate

4 Next, select **ItemActions**

5 Press **Activate**

6 Select your access level after the activation and then select **OK**

7 To return to the main menu, select **File** and then **End**.

For complete information about using SAS software on your host, refer to the SAS Companion documentation for your current operating system. For more information about the SAS Companion documentation, see "Printed Documentation" on page 67.

Using the Toolbar, Command Line, and Menus

The Toolbar
> is a group of buttons on the SAS IT Resource Management GUI for UNIX and Windows that you can select or click with your mouse in order to execute commands. *The tasks available from the toolbar are also available from the menus on the menu bar.* The items available on the toolbar are specific to the current window; therefore, the available buttons will change as you select tasks from the SAS IT Resource Management window and open new windows. To activate the appropriate toolbar for your current window, click inside the window.

The Command Box or Line
> is available in the toolbox or to the left of the toolbar, on UNIX and Windows. To execute commands from the command box, type the command in the box and press **Enter**.
>
> On z/OS, to obtain or remove a command line, issue the *COMMAND* command (typically assigned to the F10 or Shift/F10 key). To execute commands from the command line, type the command on the line and press **Enter**.

Menus
> are provided at the top of most windows in the SAS IT Resource Management GUI for UNIX and Windows, and at the top of some windows in the GUI for z/OS. Each menu contains groups of related tasks that you can perform in the window. On UNIX and Windows, many items that are available on SAS IT Resource Management menus are also available from a pop-up menu that you can activate by using the mouse. To display the pop-up menu, follow these steps:
>
> **1** Position your mouse pointer in the window where you want to perform a task or at the location in a window where you want the action to occur.
>
> **2** Click the mouse menu button (typically the right mouse button).
>
> **3** Choose a command from the list.
>
> **4** Click the mouse selection button (typically the left mouse button) to perform the selected task.

SAS Windows

When you use SAS IT Resource Management to produce graph and text reports or to run other programs, you might also want to use the following standard SAS windows to submit code, review program messages, or display output:

Log
> contains a log of your SAS IT Resource Management session and also displays anything you explicitly print to the Log window, such as table and variable definitions. This window records notes, warnings, and error messages about your job. If you are running a job in batch, on UNIX and Windows, the log file is saved by default, as *jobname*.log. On z/OS, the log file is saved with the output from your job.

Output
> is the default destination for any text reports that you produce by using SAS IT Resource Management. This window cannot display high-resolution graphics. If you are running a job in batch, by default, on UNIX and Windows, the output file is saved as *jobname*.output. On z/OS, the default output location is designated by each individual site.

Graph
> is the default destination for any graph reports that you produce by using SAS IT Resource Management. You can also direct graphics output to a SAS catalog GRSEG entry, various printing devices, or an external graphics stream file, or you can make it Web-accessible.

Program Editor
> is used to submit programs to the underlying SAS session.

C H A P T E R

3

Working with Your Performance Data

Overview

This chapter describes the following functions:

☐ How to create a PDB

☐ How to process data into the PDB

☐ How to reduce data into the summary levels of the PDB

☐ How to use the wizards within SAS IT Resource Management

☐ How to access data on the server by using a remote profile

To create a PDB and to process and reduce data, you must start SAS IT Resource Management on a host where the server software is installed.

How to Create a PDB
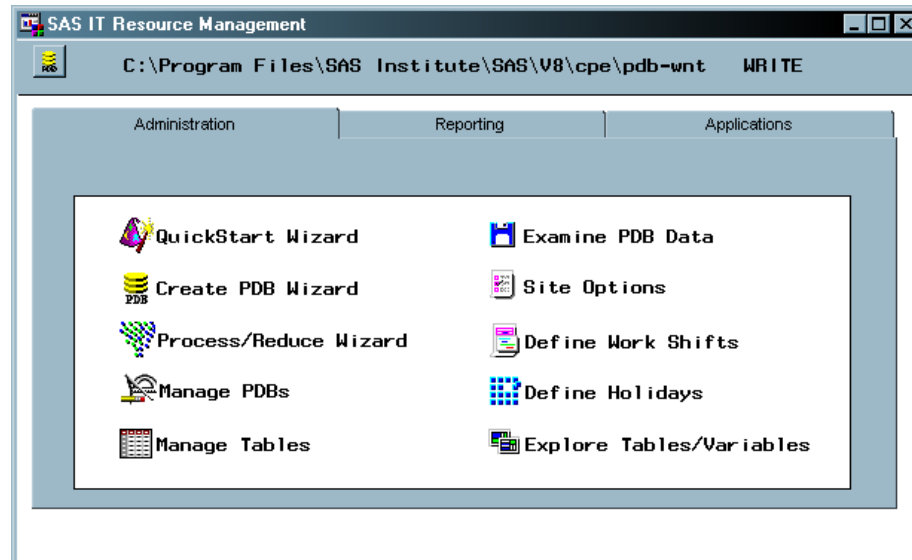
Within the SAS IT Resource Management GUIs, you can create a PDB by using the QuickStart Wizard or by using either the Create PDB Wizard or Define New PDB Wizard.

The **QuickStart Wizard** and **Create PDB Wizard** are available from the main window in the SAS IT Resource Management GUI for UNIX and Windows.

QuickStart Wizard and **Define New PDB** are first two items on the menu that opens after you select **PDB Admin** in the SAS IT Resource Management GUI for z/OS.

The QuickStart Wizard enables you to create a new PDB, to identify the source of the data that you want to process and store in the PDB, and to identify the location of your raw data file and the location for your Web reports. The wizard also creates jobs that you can run in batch, to perform these tasks:

- □ Add tables in the PDB, for your data source
- □ Process data into the PDB
- □ Reduce data into the summary levels in the PDB
- □ Create a variety of Web-enabled reports through which you can analyze the data by using a Web browser.

You can use the QuickStart Wizard to create a PDB and to create batch jobs. You can then run the batch jobs to perform the tasks listed above. On UNIX or Windows, you can also use the Create PDB Wizard to create your PDB and the Process/Reduce Wizard to process and reduce your data. On z/OS, the equivalent to the Process/Reduce Wizard is available on the menu that opens when you select **PDB Data**.

The Create PDB Wizard enables you to create a new PDB and add tables to the PDB. SAS IT Resource Management supplies tables for many popular collectors. When you create a new PDB, you can select a data category and then select the supplied tables to use in your PDB. You can use the tables as provided, or you can review the variables within the PDB and customize the tables for your site's needs. The Create PDB Wizard creates a new PDB by prompting you for the following information:

- □ the location where you want to create the new PDB
- □ the collector that you are using to collect your data

□ the tables that you want to use in your PDB

□ the options for the PDB.

On UNIX and Windows, the `Process/Reduce Wizard` guides you through the process and reduce tasks. On z/OS, the `Process Raw Data` and `Reduce PDB Data` menu items control the same tasks. These wizards are described in detail later in this chapter.

Using the QuickStart Wizard

The QuickStart Wizard provides a quick and easy way for you to begin using SAS IT Resource Management. The wizard helps you set up your data with SAS IT Resource Management by creating a new performance data warehouse (PDB), by enabling you to identify the type of data that you want to store in the PDB, and by creating SAS jobs that you can use to process and reduce data and to create reports.

The QuickStart Wizard helps you do the following tasks:

□ Create a new PDB.

□ Identify the collector that you are using to gather data.

□ Identify the category of the data that you will store in the PDB.

□ Identify the location where you want to store your Web reports.

□ Create three or more SAS jobs that you can use in batch mode. The jobs add tables to your PDB, process your performance data, reduce the data, and create reports based on that data.

□ Create the directory "under" *qs* (on UNIX or Windows) or the PDS .QS (on z/OS), that holds the SAS jobs. For more information about these locations, see "Where Are the Programs Stored?" on page 29. On z/OS, the device where new PDSs will be created is based on the FILEDEV SAS System option. (If this device is temporary, then copy the PDSs to a permanent location.)

To navigate through the QuickStart Wizard, use the **Next** and *Back* buttons in each window. To complete the QuickStart Wizard, follow these steps:

1 Start SAS IT Resource Management as described in "Starting SAS IT Resource Management" on page 14.

2 On UNIX or Windows, from the `Administration` tab in the main SAS IT Resource Management window, select `QuickStart Wizard`.

On z/OS, you can start the QuickStart Wizard by following this selection path: **PDB Admin ▶ QuickStart Wizard**.

3 From the initial window in the QuickStart Wizard, select the collector that you are using to collect the data that will be stored in this PDB. Then, select **Next** to continue.
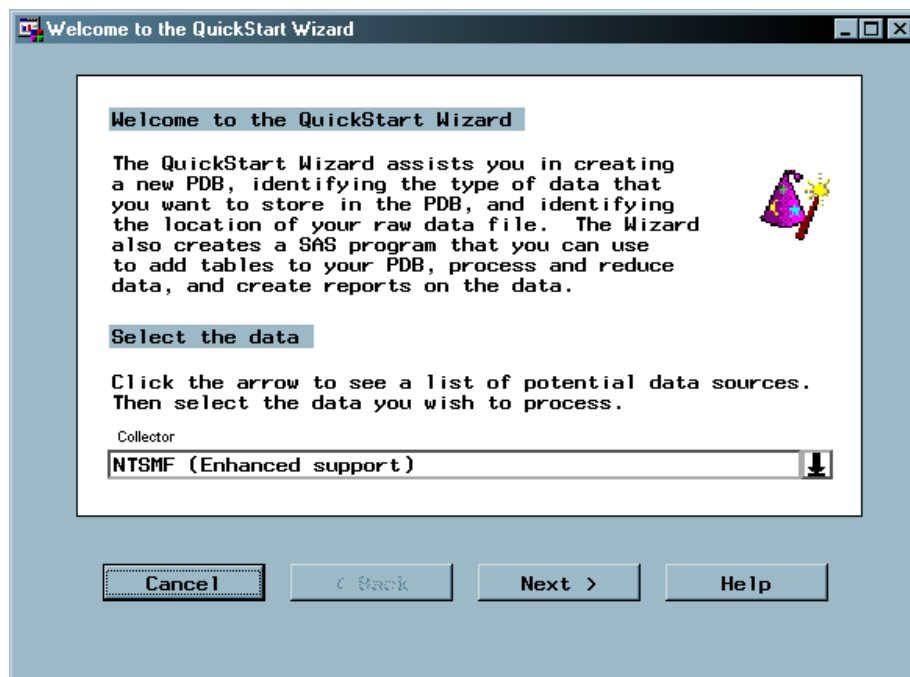
4 From the Data categories window, select the items in the list that characterize the kind of performance data you are collecting. Items that are grayed out are not currently available. The wizard will use the selected data category (or categories) to determine which tables to add to your PDB. Select **Next** to continue.
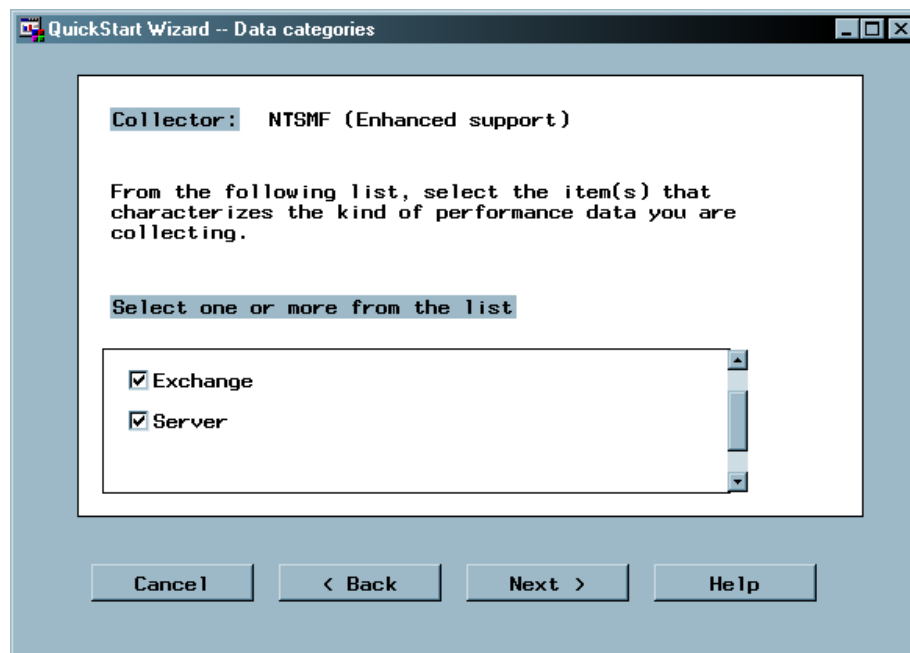
5 In the Input data location window, specify the complete name of the raw data file, including the directory path or physical location where the raw data file is stored. If this file does not currently exist, specify the location where the file will exist when it is created. If necessary, you can modify the location later by editing the processing job that is created by the wizard.

If the data is stored on UNIX or Windows, you can type the complete directory path and filename of the raw data file. Do not enclose the name in quotes. (You can also select this information from a list that can be accessed by using the arrow next to the text box.) For example, on Windows, the complete path and filename might be *c:\mydatafiles\datafile1*. On UNIX, it might be */directorypath/mydata/datafile1*. If the data is located on a z/OS host, type the fully qualified path and name, such as *my.raw.datafile*.

Click **Next** to continue.

6 In the PDB Location window, specify the complete name and location where you want the new PDB to be created. This includes the name of the PDB and the complete directory path or physical location where the PDB is stored. If the PDB already exists, you can enter the current location and name of the PDB. If you specify a PDB location that already exists and if you have modified any tables in that PDB for the data type that you have selected, the QuickStart Wizard programs that will be created by the wizard might not work with the existing PDB.

CAUTION:
Do not use a temporary file as your Data location. If your default is a temporary file or directory, your data will be deleted at the end of your session. △

Click **Next** to continue.
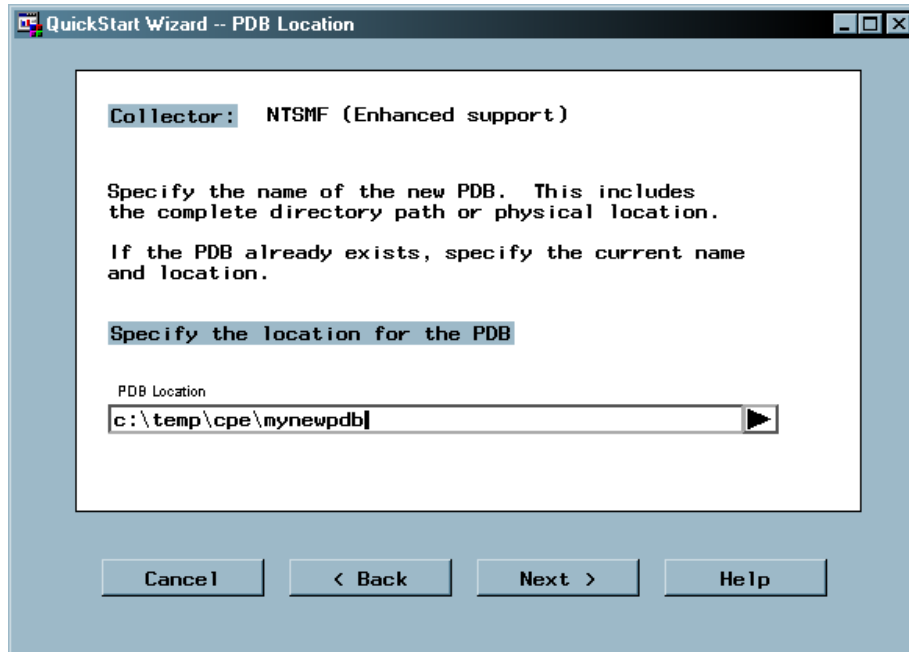
On z/OS, type the high-level qualifier, such as *sas.itsv.mypdb*. If you are creating a new PDB, you are prompted to create the PDB by completing the PDB Allocation window. If you need assistance, click **Help** in that window. Click **Run** to create the PDB and libraries. When they are successfully allocated, click **Goback** to continue with the wizard.

On UNIX or Windows, you can use the arrow beside the **PDB Location** text box in order to browse your current directory structure. If you are creating a new PDB, you can use the browse arrow to identify the path where the new PDB will be stored. Select **OK** to return to the PDB Location window, and then add the actual PDB name to the end of the path in the text box. For example, on Windows, the complete path and PDB name might be *c:\sas\itrm\mypdb* and on UNIX it might be */apps/sas/itrm/mypdb*.

CAUTION:
   Do not use a temporary file as your PDB location. If your default is a temporary file or directory, your PDB will be deleted at the end of your session. △

Select **Next** to continue.

7 You can use this window to specify where you want your Web reports to reside. You can either use the default location that is displayed on this window, or you can specify a completely different location. Your Web reports do *not* have to reside in the same location as your PDB.
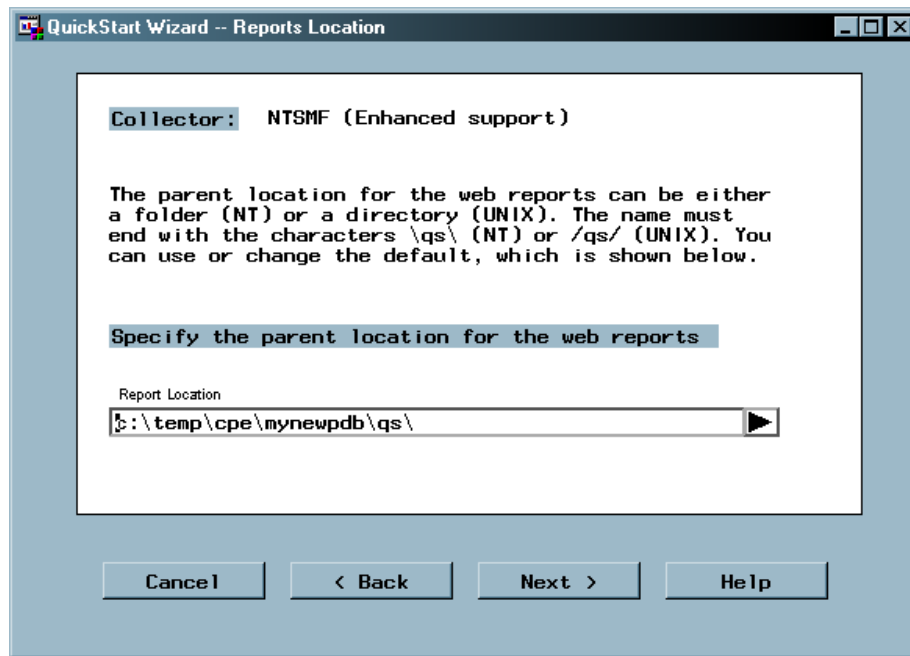
For detailed information about how to specify this location, select the **Help** at the bottom of the window.

CAUTION:

Do not use a temporary file as your *Report location*. If your default is a temporary file or directory, your report will be deleted at the end of your session.　△

Note:　On z/OS, you can direct your Web output to a UNIX File System area that is completely separate from your PDB, where you can view the reports by using your Web browser. Or you can direct your Web output to a PDS, from which you can FTP it to a host on which you can view the reports by using your Web browser. Click **Next** to continue.　△

8 On z/OS, verify that the job card information is correct. If it is not correct, edit the information in the text box. This information will be used in the SAS programs that are created by the wizard. Before you run the SAS jobs, you might need to change the default job name in each job. Select **Next** to continue.

9 The QuickStart Wizard Summary window displays the tables that must be available in your PDB when you run the process job that is created by the QuickStart Wizard. If the tables do not already exist in the PDB, they will be added when you run the QuickStart program.

If you are adding tables to an existing PDB and if any of those tables already exist in that PDB, they will not be replaced, and the *x*PROCESS job that is created in this run of the wizard will fail. This is because the QuickStart Wizard might add some variables for some tables used by the *x*PROCESS job.

```
┌─────────────────────────────────────────────────────────────────────┐
│ ⬚ QuickStart Wizard Summary                                 _ □ ✕    │
│                                                                        │
│   ┌──────────────────────────────────────────────────────────────┐   │
│   │  QuickStart Wizard Summary                                    │   │
│   │                                                                │   │
│   │  Collector:  NTSMF (Enhanced support)                         │   │
│   │                                                                │   │
│   │  Your PDB must contain the following tables for the           │   │
│   │  QuickStart Wizard programs to run properly.  If the          │   │
│   │  tables do not already exist in the PDB, they will be         │   │
│   │  added when the QuickStart Wizard programs run.               │   │
│   │                                                                │   │
│   │  ┌────────────────────────────────────────────────────────┐▲│   │
│   │  │ Table       Status       Description                    │  │   │
│   │  │ ========================================================│  │   │
│   │  │ NTCACHE     NEW          Cache                          │  │   │
│   │  │ NTCNCTN     NEW          MSExchangeMTA Connections      │  │   │
│   │  │ NTINTRF     NEW          Network Interface              │  │   │
│   │  │ NTLGDSK     NEW          LogicalDisk                    │  │   │
│   │  │ NTMMR       NEW          Memory                         │▼│   │
│   │  │ NTMSEDB     NEW          MSExchangeDB                    │  │   │
│   │  │ ◄│                                                   │►│  │   │
│   │  └────────────────────────────────────────────────────────┘  │   │
│   └──────────────────────────────────────────────────────────────┘   │
│                                                                        │
│     ┌─────────┐   ┌─────────┐   ┌──────────┐   ┌─────────┐           │
│     │ Cancel  │   │ < Back  │   │  Next >  │   │  Help   │           │
│     └─────────┘   └─────────┘   └──────────┘   └─────────┘           │
│                                                                        │
└─────────────────────────────────────────────────────────────────────┘
```

When you are ready to create your PDB and the associated QuickStart Wizard programs, select **Next** to continue.

10 The Progress Indicator window displays a status box, which you can use to follow the tasks that the wizard performs as it creates your processing jobs.

The progress indicator will display a message when the programs are complete. Select **Next** to continue.

11 The last window in the wizard displays the name of your new PDB, which is now the active PDB.

```
┌─────────────────────────────────────────────────────────────────────┐
│ ⬚ QuickStart Wizard -- Stage 1 Complete                     _ □ ✕    │
│                                                                        │
│   ┌──────────────────────────────────────────────────────────────┐   │
│   │                                                                │   │
│   │  Collector:  NTSMF (Enhanced support)                         │   │
│   │                                                                │   │
│   │                                                                │   │
│   │  Stage 1 of the QuickStart Wizard is now complete!            │   │
│   │                                                                │   │
│   │                                                                │   │
│   │  The following PDB is now active:                             │   │
│   │                                                                │   │
│   │  ┌────────────────────────────────────────────────────────┐  │   │
│   │  │ c:\temp\cpe\mynewpdb                                    │  │   │
│   │  │ ◄│                                                   │►│  │   │
│   │  └────────────────────────────────────────────────────────┘  │   │
│   │                                                                │   │
│   │                                                                │   │
│   │  Select Next to view instructions for Stage 2.               │   │
│   │                                                                │   │
│   └──────────────────────────────────────────────────────────────┘   │
│                                                                        │
│                                            ┌──────────┐               │
│                                            │  Next >  │               │
│                                            └──────────┘               │
│                                                                        │
└─────────────────────────────────────────────────────────────────────┘
```
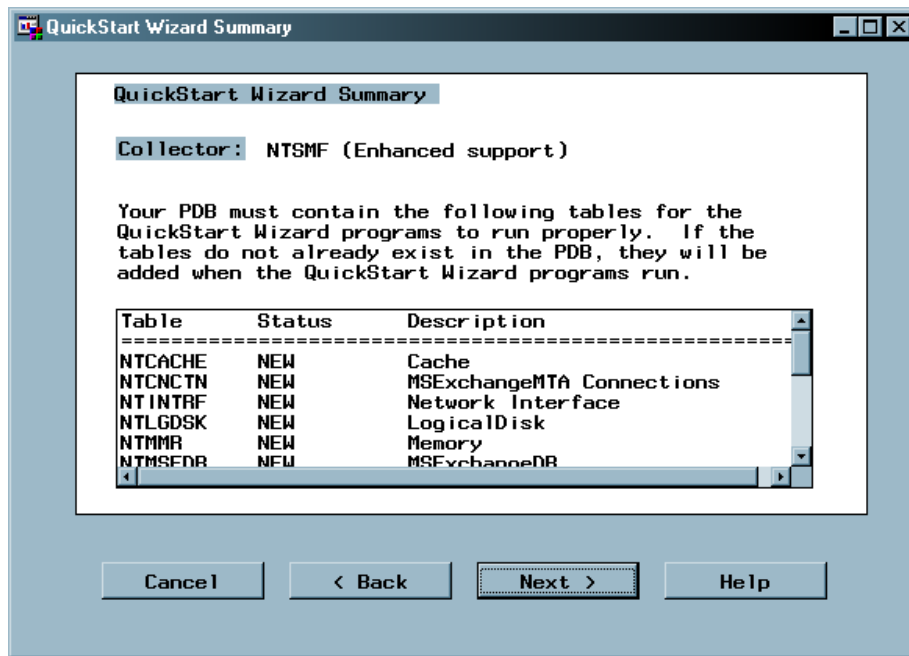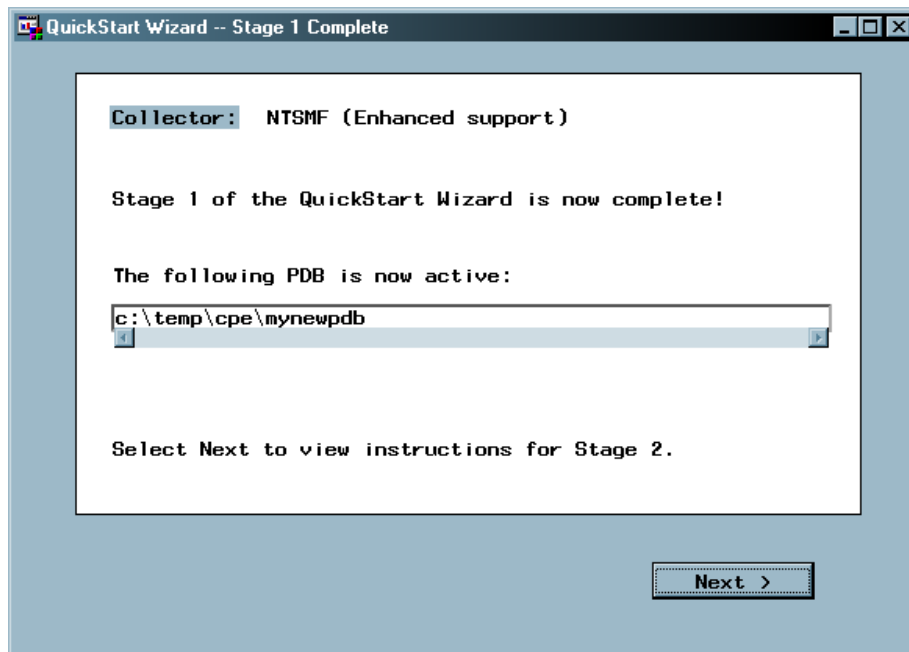
CAUTION:

Do not use a temporary file as your PDB location. If your default is a temporary file or directory, your PDB will be deleted at the end of your session. △

Click **Next** to view or print the instructions for customizing your PDB and for working with your QuickStart programs. You can also read this same information in the topic that follows, called "Customizing the PDB," and in the topic called "Customizing and Running the Programs Created by the Wizard" on page 28.

12 After reading the customization information, select **Finish** to close the wizard and to return to the main SAS IT Resource Management window or menu.

After you complete the QuickStart Wizard, perform these follow-up tasks:

☐ Customize the default PDB options and site library options for your site.

☐ Customize the *x*PROCESS and *x*REPORT jobs. On z/OS, remember to change the default jobname in the job cards.

☐ Set up your collector to collect the data that you will use with the QuickStart jobs.

Each of these tasks is described in the following sections.

Customizing the PDB

After you complete the QuickStart Wizard, the new PDB and the site library have default options. You should review the options and modify the settings to accurately reflect your site's needs.

These are the PDB and site options:

☐ PDB Shifts

☐ PDB Start of Week and Archiving

☐ Table Archiving and Age Limits

☐ Site Holidays and Holiday Shift Descriptions

You can access these tasks from the **Administration** tab in the main window on UNIX or Windows, and from the **PDB Admin** button in the main menu on z/OS.
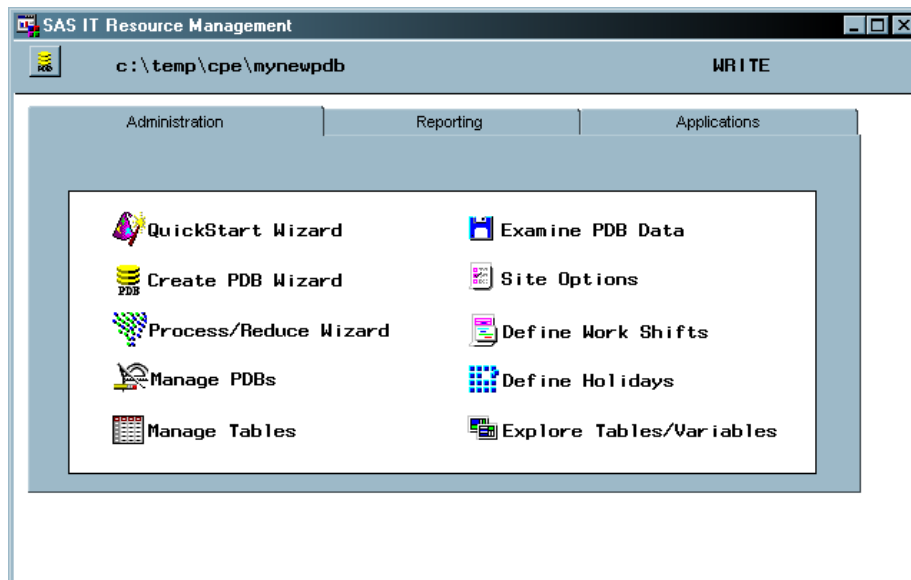
If you want to review and configure these options, follow the steps provided in the following sections. These steps enable you to access the window in which you can view

and modify the options for the PDB and your site library. For assistance in completing one of the windows, select `Help` in the window.

In the process of customizing these options, you might receive messages that indicate you should restart SAS IT Resource Management in order for the changes to take effect. If this happens, you must close SAS IT Resource Management after you have completed the customizations.

Follow this selection path on UNIX or Windows: **Exit ▶ Exit IT Resource Management Only**.

Follow this selection path on z/OS: **Exit ▶ Exit IT Resource Management only**. Then, to restart the window interface from the SAS command line, type *itrm* (or the appropriate command to start SAS IT Resource Management at your site) and press `Enter`.

Customizing PDB Shifts

On UNIX or Windows, if you want to change the PDB shifts by hour or day of the week, or to change the shift descriptions, then select the following from the `Administration` tab: **Define Work Shifts**.

On z/OS, select the following path from the main menu: **PDB Admin ▶ Define Work Shifts**.

Customizing the PDB Start of Week and Archiving

On UNIX or Windows, if you want to change the PDB start of week and archiving options, select the following path from the `Administration` tab: **Manage PDBs ▶ Properties**.

Note: This displays properties for the current PDB. △

On z/OS, select the following path from the main menu: **PDB Admin ▶ Set Active PDB Options**.

Customizing Table Archive Settings and Age Limits

On UNIX or Windows, if you want to change the archive settings and age limits for each table in the PDB, select the following path from the `Administration` tab: **Manage Tables ▶ *select a table* ▶ File ▶ Properties**.

On z/OS, select the following path from the main menu: **PDB Admin ▶ Config Active PDB Dictionary ▶ *select a table* ▶ Item Actions ▶ Edit Definition**.

Customizing Site Holidays and Shift Descriptions

On UNIX or Windows, if you want to view and change the default Site Holidays and Shift Descriptions , select the following from the `Administration` tab: **Site Options ▶ *select icon in upper left of window* ▶ Write ▶ OK ▶ Holiday or Shifts**.

On z/OS, select the following path: **Exit ▶ Exit IT Resource Management only ▶ Tools ▶ Options ▶ Command ▶ at the command prompt, type *itrm siteacc=old* ▶ OK**. Then, select the following path: **PDB Admin ▶ Define Work Shifts ▶ Site Defaults ▶ Define Holidays**.

Note: You must have write access privilege in order to modify site information. △

Customizing and Running the Programs Created by the Wizard

As you proceed through the QuickStart Wizard, you are asked to specify the location of the PDB that you want to create or update. The wizard creates several files and

directories (on UNIX and Windows) or PDSs and members (on z/OS) that are associated with your new PDB. Two of these files contain the SAS programs that you will use to process data into the PDB, reduce the data into the summary levels of the PDB, and analyze the data through reports.

Nightly, you will run the *x*PROCESS job (on z/OS) or the *x*process.sas job (on UNIX or Windows) to process and reduce data, and the *x*REPORT job (on z/OS) or the *x*report.sas job (on UNIX or Windows) to generate reports on the data. The *x*RPTSTR job (on z/OS) or the *x*rptstr.sas job (on UNIX or Windows) needs to run only one time, and the wizard runs it. (The wizard runs the *x*RPTSTR job or *x*rptstr.sas job to set up the structure for the Web reports that will be written by the *x*REPORT job or *x*report.sas job. You can modify this report structure at a later time.)

Where Are the Programs Stored?

These three programs are named and stored in directories according to your operating environment.

Table 3.1 Programs Created by the Wizard

Platform	What the programs are named	Where the programs are stored
UNIX	*x*rptstr.sas *x*process.sas *x*report.sas	in a directory called qs/cntl
Windows	*x*rptstr.sas *x*process.sas *x*report.sas	in a directory called qs\cntl
z/OS	*x*RPTSTR *x*PROCESS *x*REPORT	in a PDS called QS.CNTL

Note: *x* represents the collector or host, such as P for HP OpenView Performance Agent, N for NTSMF, or S for SAP. △

Here are examples of how program names are formed:

On Windows, if you have specified that your PDB is to be created in *c:\sas\itrm\mypdb*, then the default location where your process job will be created is *c:\sas\itrm\mypdb\qs\cntl\xprocess.sas*.

On z/OS, if you have specified that your PDB is to be created in *SAS.ITRM.MYPDB*, then the default location where your process job will be created is *SAS.ITRM.MYPDB.QS.CNTL(xPROCESS)*.

If you are running on z/OS and you do not direct the Web reports to a UNIX File System area, the wizard also creates a job named *QS.CNTL(xFTPHTML)*. The *xREPORT* job creates reports and the *xFTPHTML* job transfers the reports to a host where you can view them with a Web browser.

The following window provides a sample view of the directories that are available in a Windows Server PDB after you run the QuickStart Wizard.
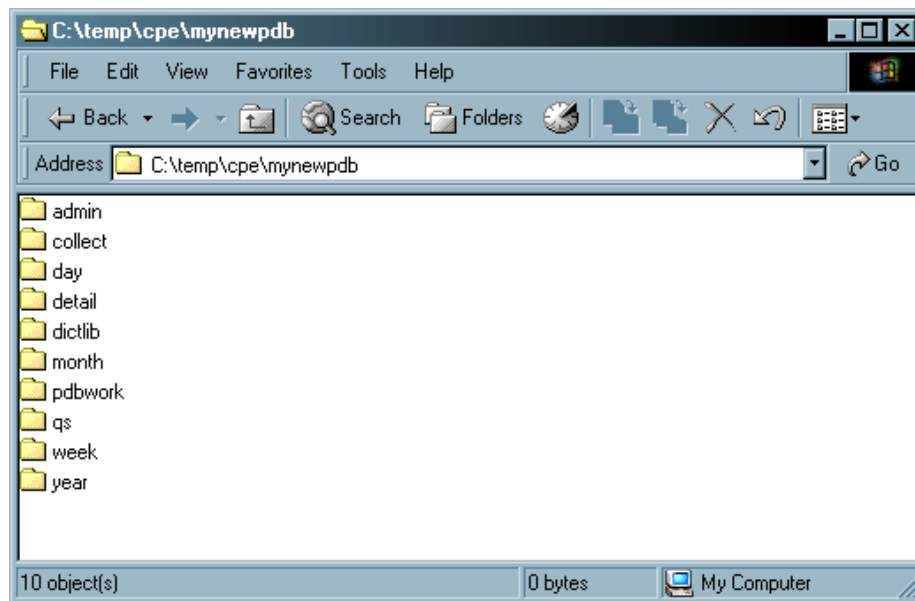
The QuickStart Wizard also creates several other *qs* directories or *.QS* PDSs, depending on the type of data to be processed.

What Tasks Are Performed by Each Program?

The xrptstr.sas or xRPTSTR job

The job contains SAS IT Resource Management macros that create the report structure in which the Web-enabled reports will be located. The wizard created and ran this program for you. The program is saved so that you can refer to it in case you want to change the structure at some later time.

The xprocess.sas or xPROCESS job

The job contains SAS IT Resource Management macros that perform tasks such as the following:

   □ Start the application in batch (by using the %CPSTART macro).

   □ Process the data from your raw data file into the detail level of your PDB (by using the appropriate %CxPROCES macro for your data).

   □ Reduce that data from the detail level of your PDB into the summary levels of the PDB (by using the %CPREDUCE macro).

You must run this job before you can run the *x*report.sas or *x*REPORT job.

The xreport.sas or xREPORT job

The job contains macros that perform the following task:

   □ Create reports on the data by using a variety of report macros. A separate group of report macros is included for each type of data that you selected in the wizard. You might want to use all of these reports or only some of them. You might also want to modify some of them for your site.

   Note: On z/OS, if you have not specified that the report structure is to be placed in the UNIX File System area, then you can transfer your report files to

another host with the commands that are available in the QS.CNTL(*x*FTPHTML) member. You can customize the FTP information in this job to perform the transfer automatically after the reports are complete. (The Reports Location window lets you specify that the report structure is to be placed in the UNIX File System area. For more information about this window, see "Using the QuickStart Wizard" on page 21.) △

You cannot run the *x*report.sas or *x*REPORT job until you run the *x*process.sas or *x*PROCESS job. Additionally, these programs are designed to run in a batch or in a background environment. You can run the programs in an interactive SAS session. To do so, you must exit from the SAS IT Resource Management GUI to SAS, and then submit the programs from the SAS Program Editor window. On z/OS, be sure to remove the JCL.

Note: If you are on z/OS, it is not recommended that you run these programs interactively on TSO. (Interactive sessions that are executed on TSO typically receive fewer system resources than batch sessions, and so interactive sessions take longer to complete than batch sessions.) △

Customizing the Programs

After you complete the QuickStart Wizard, you should review the SAS programs to make sure that they meet the needs of your site. The comments within each program contain specific instructions that will help you modify the programs.

If you are running on z/OS, you can review the job card. You might choose to modify the default jobname before you submit the programs.

Note: One consequence of having the same jobname on two different jobs is that the second job will not run until the first has completed. This can be an advantage because it ensures that the process and reduce job has completed before the report job executes. △

Using the QuickStart Programs in a Production Environment

After you customize the batch jobs created by the QuickStart Wizard, you can use a scheduling system on your host, such as *cron* on UNIX or the *Scheduled Tasks* function on the Control Panel on Windows. You can then run these programs as production jobs each night.

When you run production jobs, you will want to review your site backup procedures to ensure that you make a copy of the PDB before or after every process/reduce run. Alternatively, you might want to add additional reports to the QuickStart report program.

If you have problems running the programs created by the QuickStart Wizard, refer to the recovery and restart procedures in the *RESTART* file in the QuickStart CNTL directory or PDS.

Where to Go from Here

For more information about using the Web reports, see the topic *Directing a Report to the Web* in the Working with Report Definitions chapter of the SAS IT Resource Management User's Guide and the topic *Updating an Existing QuickStart Report* in the Working with Galleries chapter in the SAS IT Resource Management User's Guide.

For specific information about the batch macros used in these jobs, refer to the SAS IT Resource Management macro reference documentation.

For more information about how to access documentation, see Chapter 5, "Additional Documentation and Resources," on page 65.

## Setting Up Your Collector

Based on the collector and the category of the data that you specified within the wizard, an instruction file has been created to assist you in setting up your collector, so that you can collect the data that will be read by the QuickStart Wizard. The instructions are created in the QuickStart directory (qs) or PDS (.QS). The instructions are stored in a subdirectory or partitioned data set whose name is based on your collector and data category.

For example, on Windows, the instruction file is created in *qs\sccc\instruct.htm*. In this example, *sccc* represents a subdirectory. The name will change to reflect your collector and data category.

s	represents your data collector. For example, you might see the letter P for HP OpenView Performance Agent or N for NTSMF data.
cccc	four-letter abbreviation that represents the category of data that you are collecting. For example, you might see SERV for server data, CICS for CICS data, or EXCH for Microsoft Exchange data.

Therefore, if your collector is HP OpenView Performance Agent and the category is server data, then the subdirectory will be *PSERV*.

The INSTRUCT.HTM file is an HTML file that you can display by using a Web browser such as Microsoft Internet Explorer or Netscape Navigator. On z/OS, if you did not specify the UNIX File System area for reports, you can view the INSTRUCT member as text, or you can FTP the file to a UNIX or Windows host and view the file by using your Web browser.

Viewing and Working with Your QuickStart Reports

When you run the *x*report.sas or *x*REPORT job that is created by the QuickStart Wizard, the report macros in this job create analysis reports based on the data that you processed and reduced in the PDB. All reports created by using this job are designed so that they can be displayed by using a Web browser. The reports are saved as *.gif* images and *.htm* files.

When you run the *x*report.sas job on a UNIX or Windows host, your reports are stored in a QuickStart directory. For example, on Windows, your reports are created in qs\web. On UNIX, the reports are stored in *qs/web*. You can also specify a different directory on the Locate Reports window during the execution of the QuickStart Wizard.

If you run the *x*REPORT job on z/OS, your reports are stored in the location you specified in the Locate Reports window during the execution of the QuickStart Wizard. This can be an UNIX File System directory and a collection of subdirectories underneath that directory, or a collection of PDSs with a high-level qualifier of the PDB name + .QS. After you run *x*FTPHTML to do the file transfers, the reports are stored in a UNIX directory or a Windows folder.

To view your reports, use your Web browser to display the *welcome.htm* file that is created in the directory or folder named *web*. The *welcome.htm* file contains a list of all data collectors and data categories that you selected when you completed the wizard. The reports are displayed in a gallery and the report descriptions are listed on the left side of the Web page. To view the reports for a specific group of data, you can select the collector and data category for that data. When you select a description, the right side of the Web page displays thumbnail images of each report in that description group. Select one of the small images in order to display a large view of that report.

Additional information about your collector is also available by selecting **HELP** on the *welcome.htm* Web page.

Through the Manage Report Definitions window in the SAS IT Resource Management client interface, you can modify any of the report definitions for the reports created with the QuickStart Wizard report program. To edit a report definition, follow these steps:

1 Select the **Reporting** tab in the SAS IT Resource Management client GUI.

2 From the **Reporting** tab, select **Manage Report Definitions** to open the Manage Report Definitions window.

3 From the list of folders, select the folder that contains the QuickStart Wizard reports, *ADMIN.ITSVRPT*.

4 From the list of report definitions, select the definition that you want to modify.

For more information about modifying a report definition, see "The Manage Report Definitions Window" on page 51.

For more information about using Web reports, see the topic *Directing a Report to the Web* in the Working with Report Definitions chapter of the SAS IT Resource Management User's Guide and the topic *Updating an Existing QuickStart Report* in the Working with Galleries chapter in the SAS IT Resource Management User's Guide.

For more information about how to access documentation, see Chapter 5, "Additional Documentation and Resources," on page 65.

Using the Create PDB Wizard

On UNIX and Windows, the Create PDB Wizard enables you to create a new PDB and add tables to the PDB. SAS IT Resource Management supplies tables for many popular collectors. When you create a new PDB, you can select a data category and then select the supplied tables to use in your PDB. You can use the tables as provided, or you can review the variables within the PDB and customize the tables for your site's needs. After you have created the PDB, you can use the Process/Reduce Wizard to process data into the PDB and then reduce the data into the summary levels of the PDB.

The Create PDB Wizard is not available in the z/OS server interface. On z/OS, the equivalent functionality is available by selecting: **PDB Admin ▶ Define New PDB**.

To create a new PDB by using the Create PDB Wizard, follow these steps:

1 Start SAS IT Resource Management, and from the **Administration** tab in the main window, select **Create PDB Wizard**.

2 From the first window in the PDB Wizard, type (or use the arrow next to the text box) to indicate the complete physical location and name of the PDB that you want to create.
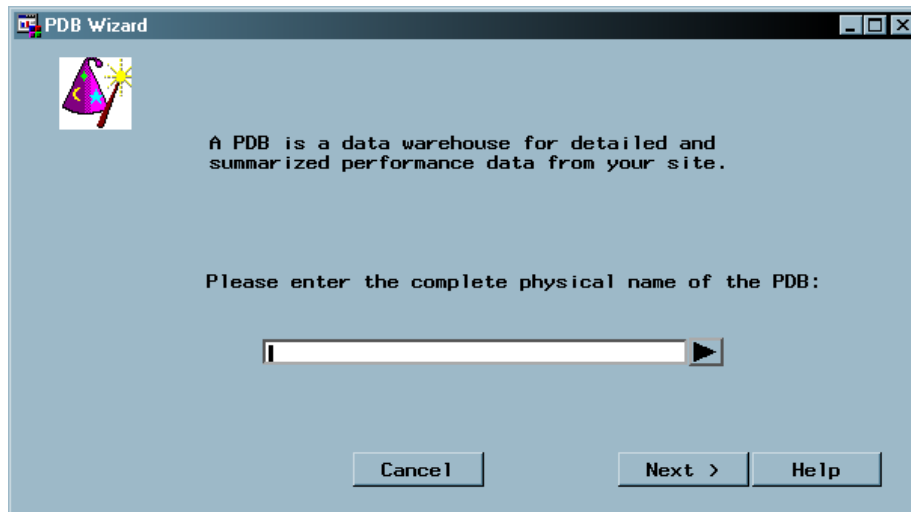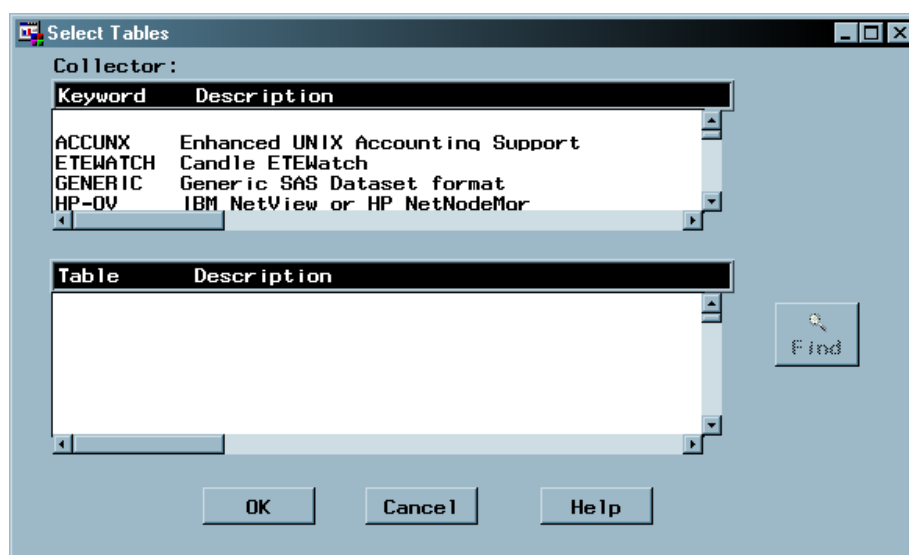
Select **Next** to continue.

3 SAS IT Resource Management provides tables for many popular data collectors. To see the list of supplied tables, choose **Select Tables**.
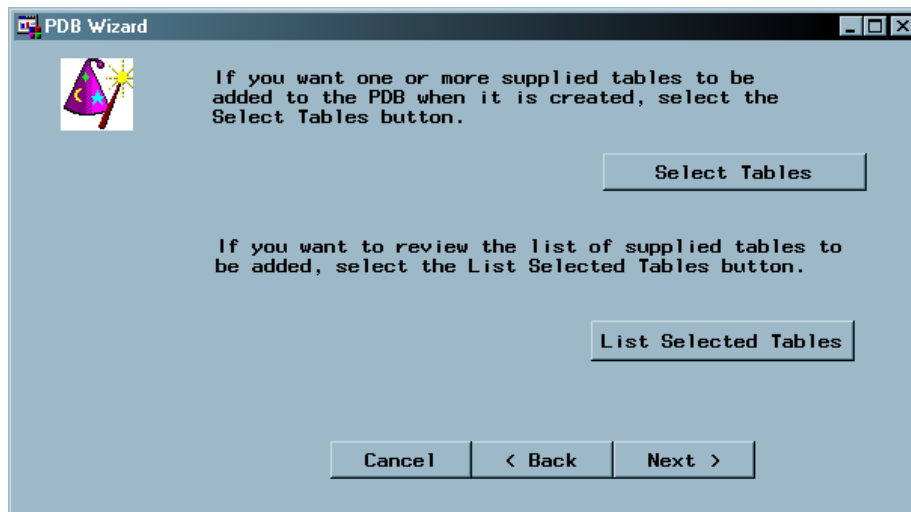
From the Select Tables window, select the collector that you will use to collect raw data.

Then, from its list of tables, select the tables that you want to add to your PDB. Select **OK** to close the window.

When you are finished selecting the tables to add to the PDB, select **List Selected Tables** to review the list of tables. When you are satisfied with the list, select **Next** to continue.

4 Select the day of the week to use as the start of week for the new PDB. Select **Next** to continue.

5 Indicate whether you want the new PDB to be your active PDB when you have finished working with the wizard. Select **Next** to continue.

6 Review your PDB selections on the last window of the wizard, and when you are satisfied with your choices, select **Finish**. This will create the PDB and close the wizard.

After you create your PDB, you should review and customize the following default PDB options and site library options:

□ Work Shifts

□ Site Holidays and Shift Descriptions

□ Start of Week and PDB Archive Status

□ Table Archive Status (for each table in the PDB)

On UNIX or Windows, these options are available from the **Administration** tab in the main SAS IT Resource Management window. To view these options, select the following path:

□ **Define Work Shifts**

□ **Site Options**

□ **Manage PDBs**; then, select **Properties**

□ **Manage Tables** and select a table from the list of tables that is displayed. Then, select **Properties** from the pull-down **File** menu. Repeat for each table.

On z/OS, these options are available from the **PDB Admin** button in the main SAS IT Resource Management menu. To view these options, select the following paths:

□ **Define Work Shifts**

□ **Set Active PDB options**, and

   □ **Config Active PDB Dictionary**, and select a table from the list that is displayed. Select **Item Actions** and then, select **Edit Definition**.

You can access, view, and run reports on the data in the PDB from either the client or the server. However, you cannot modify the PDB in any way without invoking SAS IT Resource Management on the server.

When you have set up your PDB and your collector you can process and reduce data into the PDB.

Using the Process/Reduce Wizard

The Process/Reduce Wizard enables you to process or load raw data into an existing Performance Data Warehouse (PDB) and to reduce data into summary levels in the PDB. The Process/Reduce Wizard is not available on the z/OS server platform. On z/OS, the equivalent is available by selecting **PDB Data**, and then **Process Raw Data** and, later, **Reduce PDB Data**. However, it is much faster on z/OS to do the process task and the reduce task in batch mode. (Interactive sessions that are executed on TSO typically receive fewer system resources and so they take longer to complete.)

Note: When you complete each wizard, you have the option to process and reduce your data at that time or to save the macro source code in a file. If you save it, you can then run this saved macro source code in batch. If you process and reduce the data interactively, be aware that it can take a considerable amount of time, especially if you are processing a large amount of data. The speed of your network and the power of your local host can also affect your processing time. △

Processing Data into the PDB

You can use the Process/Reduce Wizard to process data into one or more tables in your PDB. After you have processed the data, you can choose to continue with the wizard and also reduce the data, or you can choose to stop the wizard and reduce the data at a later time.

To process data into the PDB, follow these steps. (Use the **Next** and **Back** buttons in each window to navigate through the wizard.)

**1** Select the **Process/Reduce Wizard** item on the **Administration** tab.
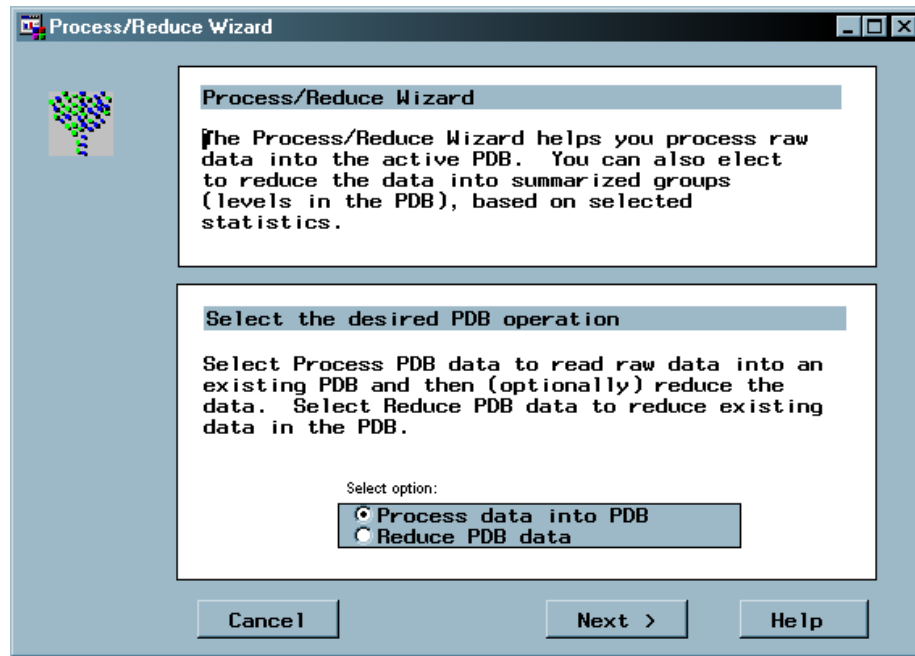
2 From the Process/Reduce window, select `Process Data into the PDB`. Select `Next` to continue.

3 This window opens.
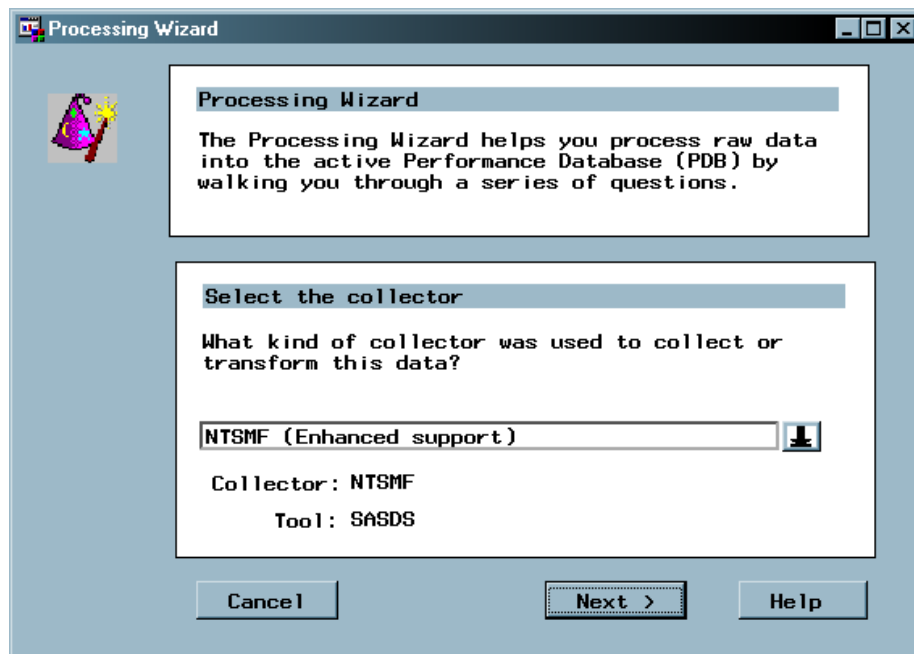
Then, select the name of the collector that was used to collect the data that you want to process into the PDB. Select `Next` to continue.

4 A window opens and enables you to specify the tables you want to use.
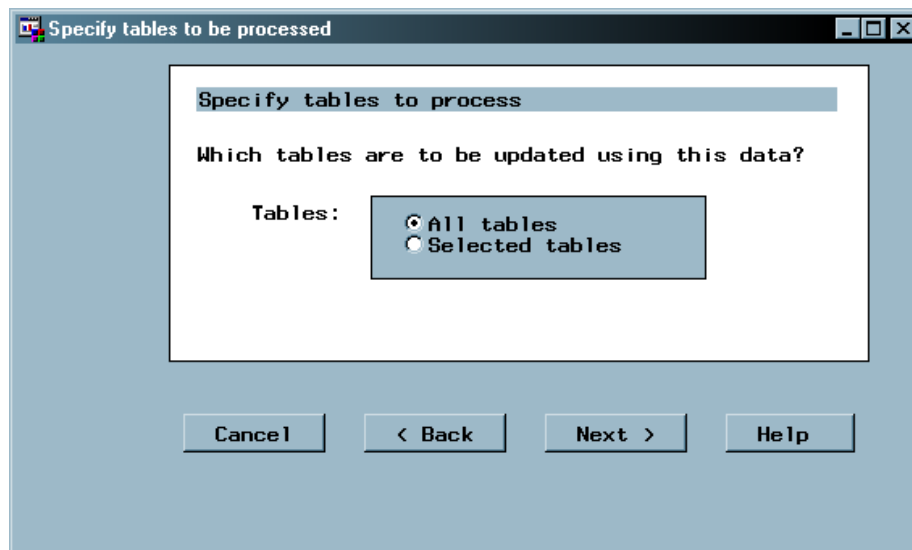
Indicate the tables into which the data will be processed. Select **All** to process data into all tables whose Kept property is **Yes**. To process data into a selected list of tables, choose **Selected Tables** and then select the tables from the list. Select **Next** to continue.

5 Now, you can specify the location of the raw data that you want to use as input.
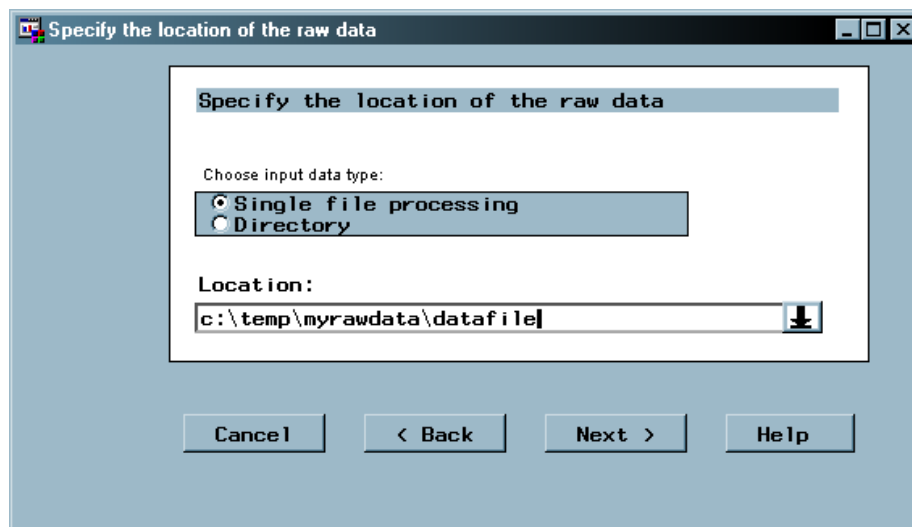
Type the complete directory path and filename or the complete physical location of the raw data file that you want to process. You can type the name in the text box or select the file by using the down arrow next to the **Location** field. Select **Next** to continue.

6 On the next few screens, you will be prompted to provide specific information about your data, based on the collector that you specified. If you want help to complete these screens, you can find additional collector setup information online within SAS IT Resource Management Server Setup Guide. For information about the locations of the Server Setup Guide, see Chapter 5, "Additional Documentation and Resources," on page 65.

7 Depending on the collector that you have selected, you have the opportunity to set or modify several advanced options that affect how your data is processed into the PDB. If you do not set these options, default values will be used.

To select or modify these options, select **Advanced**. To use the default options, select **Next**.

Advanced options include:

- **Age Limits** – This option determines how long to retain the data in a specified level of the PDB. If the data is older than the age limit, the data is aged out or removed from the PDB. (See also "How Long Is Data Stored in the PDB?" on page 8.)

- **Date Checking** – This option indicates the range of data to be processed into the PDB. Only data that falls within that range will be processed into the PDB.

- **Subsetting Data** – This option specifies a subsetting expression that will be used to determine which data is processed. Only data that satisfies the expression will be processed into the PDB.

- **Log Format** – This option indicates whether the data log will contain metrics from one table (normal) or multiple tables (multitab). (This option only applies to certain collectors.)

- **Optimization Options** – These options enable you to optimize the task of reading and processing the data based on CPU time (the time it actually takes to process the data), or based on temporary disk space (the amount of disk space that the data will use while it is being processed).

- **Macro return code** – This option specifies the name of a return code variable. For example, when your data has been processed, this variable will return the status. If errors occur during processing, the value is returned through this variable. (For more information about return codes, refer to the %CPFAILIF macro or the %CPDEACT macro in the SAS IT Resource Management macro reference documentation.)

- **User exits** – This option indicates the location of your exit source code, if you are utilizing exits during the processing step. If you are using one or more exits, identify how the exits are stored, in a SAS catalog or a directory, and then specify the complete location of the exit source code. For more information about using exits during processing, refer to the Extensions chapter in the SAS IT Resource Management User's Guide.

- **Duplicate filtering** – You can either allow or prevent duplicate data. For more information about duplicate data filtering, refer to Appendix 4 in the SAS IT Resource Management macro reference documentation.

8 From the Process Specifications Complete window, you can select **Process** to process the data now, or you can select **Save** to save the macro source code for your processing selections so that you can run the macro code at a later time or in batch.

9 After you have processed the data into the PDB, you can reduce the data. To do this, select **Close**, and then select **Continue to Reduction Wizard**. Alternatively, you can end the wizard by selecting **Return to Main Window**. To continue with the Reduction Wizard, see the instructions in "Reducing Data into the Summary Levels of the PDB" on page 39.

Reducing Data into the Summary Levels of the PDB

The reduce task within the Process/Reduce Wizard enables you to reduce data from the detail level of the PDB into the summary levels in the PDB. From these summary

levels, you can easily report on your data by using the SAS IT Resource Management reporting tools.

To use the Process/Reduce Wizard to reduce data that has already been processed into the PDB, follow these steps. (Use the **Next** and **Back** buttons in each window to navigate through the wizard.)

1 Select the **Process/Reduce Wizard** item on the **Administration** tab in the window interface, or if you have just completed the Process Wizard, continue with step 3.

2 From the Process/Reduce window, select the **Reduce PDB Data** option. Select **Next** to continue.

3 Select the tables whose data you want reduced from the detail level of the PDB into the summary levels of the PDB (day, week, month, year levels). You can select **All tables** to reduce the data in all PDB tables whose Kept property is *Yes* or **Selected tables** to select the tables whose data you want to reduce. Select **Next** to continue.
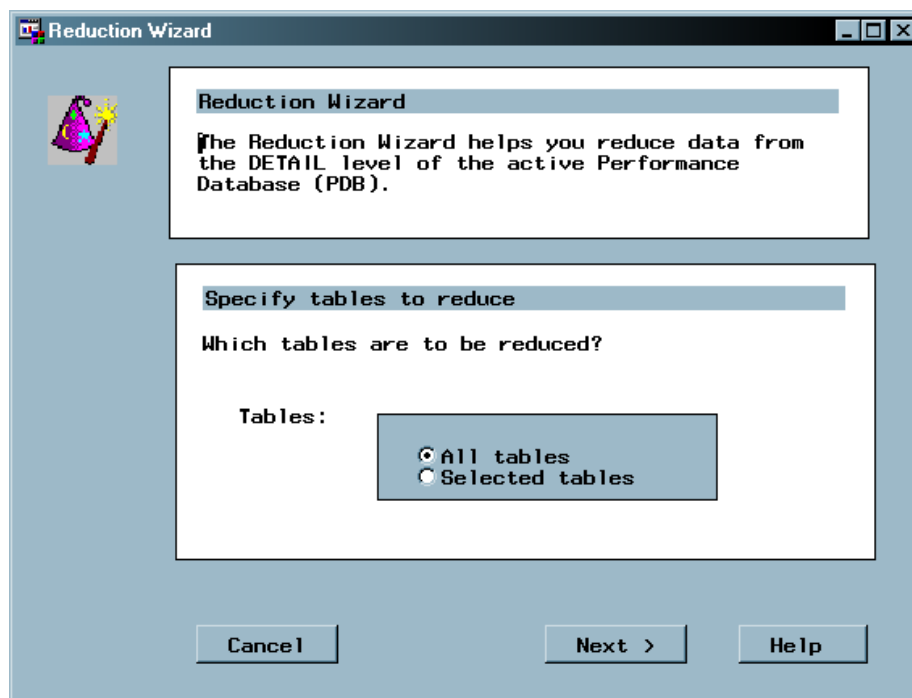
4 In the final window of the Reduction Wizard, select **Reduce** to reduce the data. From this window you can also save the macro source code to run later or to run as a batch job. When you have completed the wizard, select **Close** to close the wizard.

Now that your data has been reduced and summarized into the day, week, month, and year levels in your PDB, you can create and run reports to analyze your IT data. You might want to perform your reporting and analysis from the SAS IT Resource Management client. If your PDB is stored on a remote server host and if you will report on the data on a client host, you will need to make sure that your client is set up to connect to the PDB on the remote host. For more information, see "Analyzing Data on a Remote Server" on page 41.

Analyzing Data on a Remote Server

From the host where the SAS IT Resource Management client is installed you can access, view, and run reports by using the data in a PDB that is stored on the server. However, from the client, you have read-only access to the PDB and limited write access to other related items. To modify the PDB, you must access the PDB on the host that has the SAS IT Resource Management server license and the PDB are located.

To access a PDB that does not reside on your local host, you must create and assign a Remote Server Connection Profile for the remote PDB. The remote server profile is created on the client and then it is associated with a PDB through the Manage PDBs window in the client interface.

The profile identifies the host that has the server license and the PDB, and also defines other options such as the communications method that you will use to access the server. The profile enables you to access and read data that is in the PDB on the remote server.

To create and assign a profile to a PDB, perform these tasks:

1 Select the **Administration** tab in the main SAS IT Resource Management window.

2 Select **Manage PDBs** on the **Administration** tab.

3 From the **Locals** menu, select **Remote Profile**.

4 The Remote Server Connection Profiles window lists the available profiles. Several sample profiles are provided with SAS IT Resource Management. To use one of the sample profiles, you must modify the profile with the appropriate settings for your site.
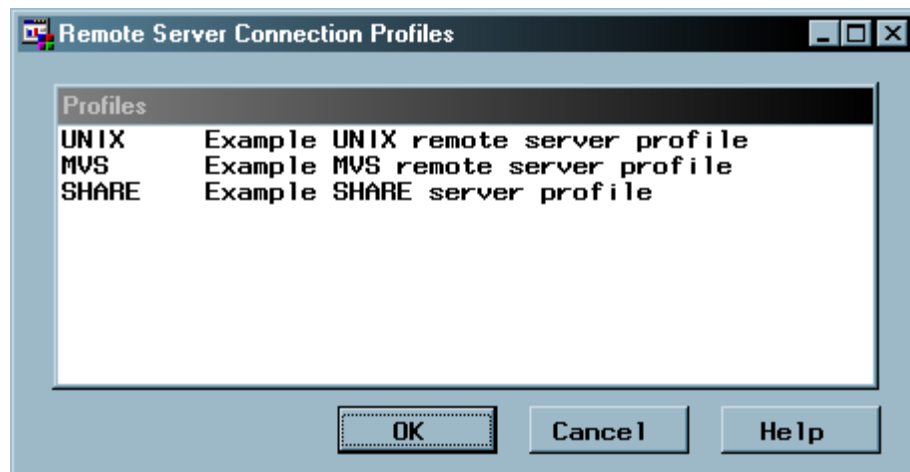
To delete a profile, select the profile from the list that is displayed. Then, select the following path: **File ▶ Delete**.

Then, select **OK**.

An easy way to create a new profile is to copy an existing profile, rename it, and then modify it.

To modify a profile, select the profile from the list that is displayed. Then, select the following path: **File ▶ Open**. The Remote Server Profile Definition window opens.
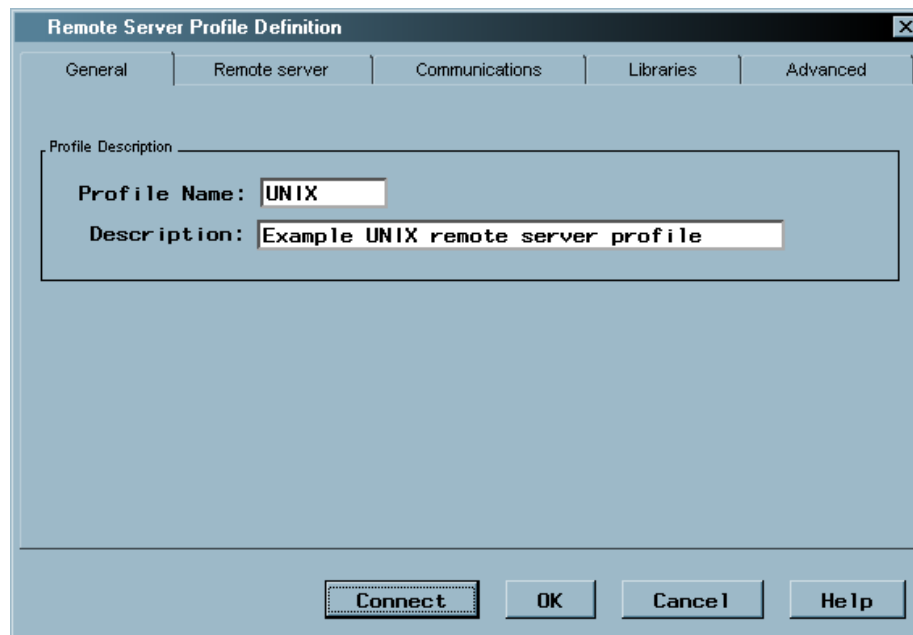
This window displays five tabs: **General**, **Remote server**, **Communications**, **Libraries**, and **Advanced**.

5 To create or copy a profile, complete the Remote Server Profile Definition window.

The following fields are available from among the five tabs on this window:

Profile Name
This field contains a name that you choose to identify this profile. It is required.

Description
This field contains a description of the profile.

Server Name
This field contains the name of the remote server. It is required. For TCP, this is an internet address. For communications methods such as APPC, this would be a logical unit name (or "luname") or a network name.

Server OS
This field specifies the operating system used on the remote server. It is required.

Server Type
This field specifies the type of server connection. You can select a single-user (CONNECT) or multi-user (SHARE) server. The default Server Type is Connect. For you to be able to connect by using SAS/SHARE, your SAS IT Resource Management server administrator must set up a SAS/SHARE server on the remote host. See your administrator to obtain details, such as the Service Name.

Note: There is also a field for indicating whether the remote server is using a spawner. △

Communications Method
This field specifies the method to be used to communicate with the remote server, such as TCP or other network access methods. The default is TCP.

Logon Script
This field specifies the name of the logon script. Some communications methods, such as TCP, require a script. Sample logon scripts are included with SAS IT Resource Management.

Note: There is also a field for specifying the name of the command that invokes SAS on the remote server. △

Location of Sitelib on server
This field specifies whether you will use the site library on the local host or on the remote host. The remote site library might contain site options specified by your server administrator.

Location of Pgmlib on server
This field contains the location of the SAS IT Resource Management program library on the remote server. If you have installed this application in the default install location, you can leave this field blank. Otherwise, it is required.

Location of ITRM Macros on server
This field contains the location of the SAS IT Resource Management macros on the remote server. If you have installed this application in the default install location, you can leave this field blank. Otherwise, it is required.

Location of MXG Format Library on server
If there is an MXG format library, this field contains the library's location.

SAS Code to submit locally and **SAS Code to submit to server**
These fields specify any additional SAS code to be submitted either locally or, for single-user (CONNECT) access, to the remote server at the time you sign on to the remote server.

Note: Depending on your communication method, some profiles might require a script. For example, TCP requires a script that is specific for your site. Sample scripts, such as *tcpmvs.scr* or *tcpunix.scr*, which can be used with many platforms (z/OS, UNIX, and Windows), are provided in *sasmisc*, in the location where SAS IT Resource Management is installed at your site. △

6 Select **OK** to close the window and create the profile or save your modifications to an existing profile.

Now that you have created the remote server profile, you must assign the profile to your remote PDB in order to access the PDB from the client. To assign the remote profile to the PDB, follow these steps:

1 Select the **Administration** tab in the main SAS IT Resource Management window.

2 Select **Manage PDBs** on the **Administration** tab.

3 Select **Add PDB Entry** from the **File** pull-down menu. Specify the PDB's name and location (but not the server).

4 Select the rightmost arrow that is associated with the **Remote Server profile** field.

5 From the Remote Server Connection Profiles window, select the profile to assign to the selected PDB.

6 Select **OK** to close the window and assign the profile.

Using SAS IT Resource Management Macros

SAS IT Resource Management provides macros that you can run in the SAS Program Editor window (outside of the SAS IT Resource Management GUI) or in batch mode. The SAS IT Resource Management GUI uses these macros. Therefore, when you select tasks within the GUI, the macros are executed to perform the tasks. If you do not want to perform tasks by using the GUI, you can write batch jobs that call these macros, and you can then submit the jobs to run in batch mode. In many cases, for example in the Manage Report Definitions window, you can make selections. Then, you can save the macro calls for those selections and run them in batch mode.

Macros that require write access to the PDB, such as %C*x*PROCES, %CPREDUCE, or any macros that update the data dictionary, must run on the SAS IT Resource Management server. Reporting macros can run on the SAS IT Resource Management client or server.

For more information about macros, see Chapter 5, "Additional Documentation and Resources," on page 65.

CHAPTER

4
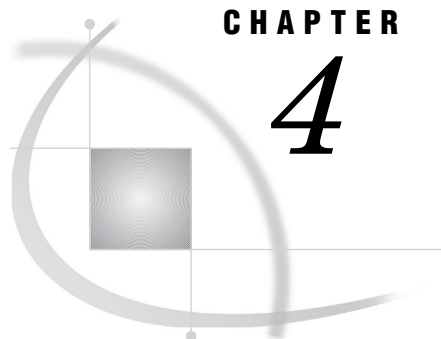
Analyzing Data through Reports

Overview

This chapter provides an overview of the reporting tools available in SAS IT Resource Management. Typically, interactive reporting is performed on the client through the GUI, and batch reporting is performed on the server by using calls to the SAS IT Resource Management macros. This chapter discusses the interactive reporting methods available within the GUI. These reporting tools are available from the `Reporting` tab in the main SAS IT Resource Management window on the client.

You can create reports on the SAS IT Resource Management server, but some interactive reporting tools that are available on the client are not available on the server.

Becoming Familiar with the Tools

The SAS IT Resource Management client GUI contains a variety of reporting tools that enable you to explore and analyze your IT resources easily and effectively.
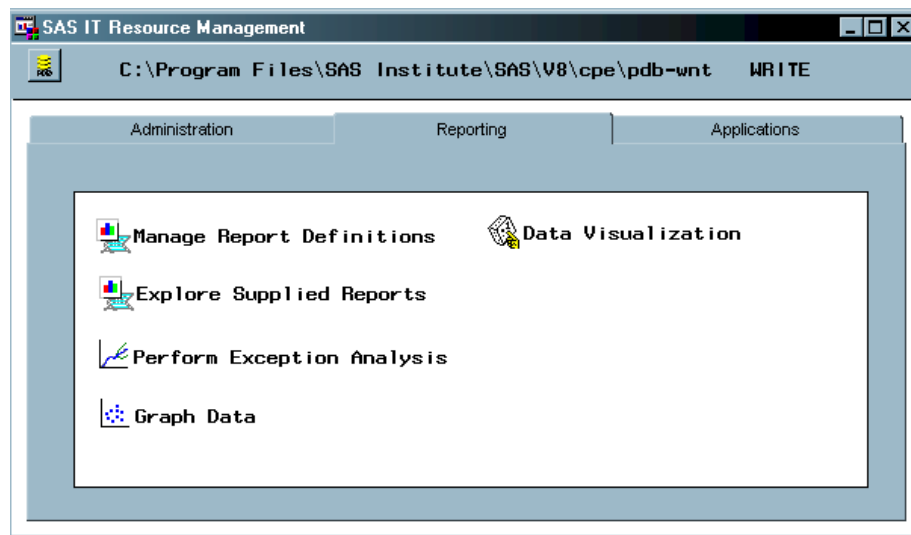
There are report tools that enable you to perform ad hoc analysis and to drill down and explore the detailed information contained in a report. These tools also help you to create exception reports that identify trends in your data, and to create high quality graphs that use patterns, colors, and symbols to represent different variables and trends in your data. Other tools enable you to create reports that you can run, save, and rerun against your production data, interactively or in batch mode.

After you have selected the report tool, you can select from a variety of report styles. You can create basic analysis reports with one or more variables, reports that group your analysis variables based on a specific variable, and spectrum reports that utilize color to display patterns in your data.

When your data has been processed and reduced into the PDB, you can analyze your data on the SAS IT Resource Management client by using the reporting tools that best meet your needs. The following table lists potential analysis reports that you might need, and the corresponding SAS IT Resource Management report tool that best meets each need.

Table 4.1   Table of Reporting Tools

If you need...	Consider using...
Production report definitions that you can run interactively or in batch mode	Manage Report Definitions
Custom report definitions that you can design, run, and save	
Reports to distribute via the Web	
Broad, ad hoc data exploration using a variety of existing report styles	Graph Data or Data Visualization
Rule definitions that identify error or exception conditions	Perform Exception Analysis

If you need...	Consider using...
Reports that you want to create and run in batch mode	The saved macro source code from Manage Report Definitions or
	The report macros in batch mode

You can access these tools from the **Reporting tab** in the SAS IT Resource Management client GUI. Start SAS IT Resource Management, and then select the **Reporting** tab from the main application window.

Many of the tools available in the **Reporting** tab (for example, Graph Data, Data Visualization, and Perform Exception Analysis) enable you to select the report style, data, colors, and other report attributes dynamically. When you create reports by using these tools, you can usually save or print the report that you create. That report is the actual output that you create by using the report tool. If you want to re-create the report, you must open the reporting tool and re-select the options to create the report.

Other report tools, such as the Manage Report Definitions window and the batch report macros, are more suited for production reports because they enable you to create and save a report definition that you can run on a regular basis against your production data. The actual report is created when you create and run the report definition against data in your PDB.

A *report definition* identifies information such as which data to use for your report, what variables to calculate and display on the report, which type of report to produce (for example, a bar or pie chart or a plot), and many other report attributes.

A *rule definition* identifies information such as which data to use for your reports, the exception conditions, and the message to generate if an exception if found.

For a closer look at each report tool, its features and advantages, and the options for running and saving reports, refer to "A Closer Look at the Report Tools" on page 51.

The Report Process

Whether you use one of the tools within the window interface or use the batch report macros, the process of creating a report involves the following steps:

1 Decide which report method or tool to use for your report.

2 Check with your SAS IT Resource Management site administrator and determine where the data for your report is located. (You will need to find out the name of the server where the PDB is stored, the name of the PDB from which the report is created, and the level and table in that PDB).

3 Create the report by using one of the tools on the **Reporting** tab in the window interface.

If you are using the Manage Report Definitions window or the batch macros, select a report style and the data to use for the report. You can then customize the report by selecting other attributes. All of these selections are used to create a report definition, which you must run to create the report.

For all other report methods, the report is created after you identify the data and the variables that you want to use for the report. You also need to identify other initial options. (These depend on the reporting tool that you use.) You can then dynamically modify the report by selecting other attributes as you analyze the data.

4 To select where and how the report will display, you must set output options for your report.

5 If you are creating the report definition by using the Manage Report Definitions window, you can generate the report by selecting **Run** in the Manage Report Definitions window.

6 If you are using the Manage Report Definitions window, you can save the report, and/or the interactive and batch forms of the report definition. The batch form can then be run (in batch mode) to create the report.

If you are using one of the other reporting options, you can print or save the report that you created.

Creating a New Report

When you create a new report definition, regardless of the reporting tool that you choose to use, you will need to select report attributes, such as the data to use in the report, the report style, and more. How you select or identify this information will vary depending on the reporting tool. You might also specify more items for one report tool than another, but the type of information that you select will be basically the same.

For example, you will need to identify the data table to use in your report, regardless of the report tool that you are using. All of the tools will use the active PDB, but some of the tools allow you to change PDBs within the report window and other tools will not let you change PDBs.

To create a report, you must at least define or select the following items:

☐ the location of the data

☐ the report style

☐ the folder for the report or rule definition

☐ the selected data to use for the report

☐ the output format and location.

Identify the Location of the Data

When you create a report, you must identify the location of the data that you want to analyze in the report. This includes the location and name of the PDB that contains the data that you want to analyze. You should verify that the currently active PDB contains the data that you want to use for the report.

For some of the report tools, when you select the PDB name, you will also need to select the table and level of the PDB. However, in the Manage Report Definitions window, your data selection occurs in a separate step. (See "Select the Data for the Report" on page 50.)

Select the Report Style

You can select from a variety of report styles, including 3-D graphs, bar and pie charts, plots, spectrum plots, and print or tabular reports. The way that you select the report style varies across report tools. For example, you can select the style from a set of viewer selection icons that run down the left side of the window in the Graph-N-Go application (Graph Data) or from the menu (Data Visualization).

Select the Folder

On the SAS IT Resource Management client, reports, report definitions, and other related items such as exception rules that are used to run exception reports are stored in folders. When you save a report or rule definition, you must identify the folder in which to store the item. When you want to run an existing report or rule definition, you must select the folder that contains that definition, and then you can select the report or rule definition that you want to run.

A list of your current reporting folders is available in the Manage Report Definitions window. You can add new folders to the list, create new folders and add them to your list, and you can remove folders from the list. For example, SAS IT Resource Management provides many sample report definitions in the *PGMLIB.ITSVRPT* folder that is shipped with this product. Similarly, a list of your current rule folders is available in **Perform Exception Analysis**.

By default, the report definitions that you see in the **Current PDB Reports** list is based on your current folder and the active PDB. The list will only contain the report definitions in the current folder that use tables that are available in the active PDB.

Select the Variables for the Report

In addition to selecting the variables that you want to display on your report, you can also assign a use to these variables. For example, when you process the data into the PDB you can elect to store many of the variables from your data, such as DISKID, MACHINE, DATE, HOUR, SHIFT, DATETIME, and so on. When you decide to create a report, you can select which variable you want to use as the BY variable, CLASS variable, GROUP variable, ANALYSIS variable(s), and so on. For example, if you want to analyze specific information about each machine, you could select MACHINE as the BY variable.

Below are some of the common variable uses in many of the SAS IT Resource Management reporting tools.

ACROSS
> This specifies the variables to use for column headings for tabular report types.

ANALYSIS
> This specifies the variables that you want to analyze in your report. If you select statistics for your report, those are applied to these analysis variable. If you specify multiple analysis variables, the unit of measure for all of the variables should be the same, and the range of values should be similar.

BY
> specifies the variables in the order that you want them sorted for your report. A separate graph or chart is produced in the report, for each value of the BY variable or for each unique combination of BY variable values, if you have multiple BY variables. For example, if you have four DISKID values on each of three machines and you select the variables MACHINE and DISKID as BY variables, you will receive 12 reports, one for each unique combination of the two variables.

CLASS
> This specifies a group or category variable on which calculations are performed. For each unique value of the class variable, a separate portion is created on the graph. For example, a portion might represent a portion of pie on a pie chart, but all values of the class variable are included in a single graph in the report.

GROUP
> This is the name of the variable that can be used to subset the observations in a chart. A group variable is useful when you want to display multiple analysis variables on one chart. The analysis variables display in groups, with one group for each value of the GROUP variable. For example, you might want the chart to display one bar for each value of the group variable that is called MACHINE.

STACK
> This is the name of the variable by which your data is subdivided. A separate bar or block is produced for each value of the STACK variable. Each of these bars (for a bar chart) is divided into sections, with one section for each analysis variable.

STATISTICS
> This specifies the statistics to use to summarize data. The available statistics vary depending on your report type.

SUBGROUP
> This specifies a variable by which to subdivide the groups of variables specified in the GROUP variable. Each subdivision of the GROUP variable becomes a portion of the graph, such as a section on a bar.

X VARIABLE
> This specifies the variable to display on the X axis for report styles that use an X axis.

WEIGHT
> This is the variable by which to weight statistical calculations for Interval data. The default weight variable is DURATION (for Interval tables, but not Event tables).

Note: Some report tools or report styles do not use all of these variable types. Variable types that are not supported for a particular report style, are not listed for that report style, and, if previously selected, are ignored. *For this reason, you should select the report style before selecting the variables and variable types.* △

Select the Data for the Report

To create a report, you must identify the data in the active PDB that you want to use for the report. This includes the table that contains the data, as well as the level of the table to use for the report. For example, you could create a report for the day level of your table, the week level, and so on. If you are creating a report by using the Manage Report Definitions window, you can also further summarize the data that is used in your report. This is done by selecting a Summary Time Period.

When you select a Summary Time Period, such as 15 minutes, one hour, or one week, your data is combined and grouped into one observation for every 15 minutes, hour, or week, depending on the time period that you select. This enables you to further summarize data at any level of the PDB. (However, you cannot see granularity finer than the level of the table. For example, you cannot see 15 minute periods in the day level, nor individual days in the week level.)

Select the Report Output Type and Location

When you create and run a report, you can select the format in which to save or print the output. For example, you might want to direct the output to a SAS window or to the Web, or print it in a specific format to a file. The format in which you can save the report output varies for each report tool.

Some common output formats are as follows:

□ SAS catalog

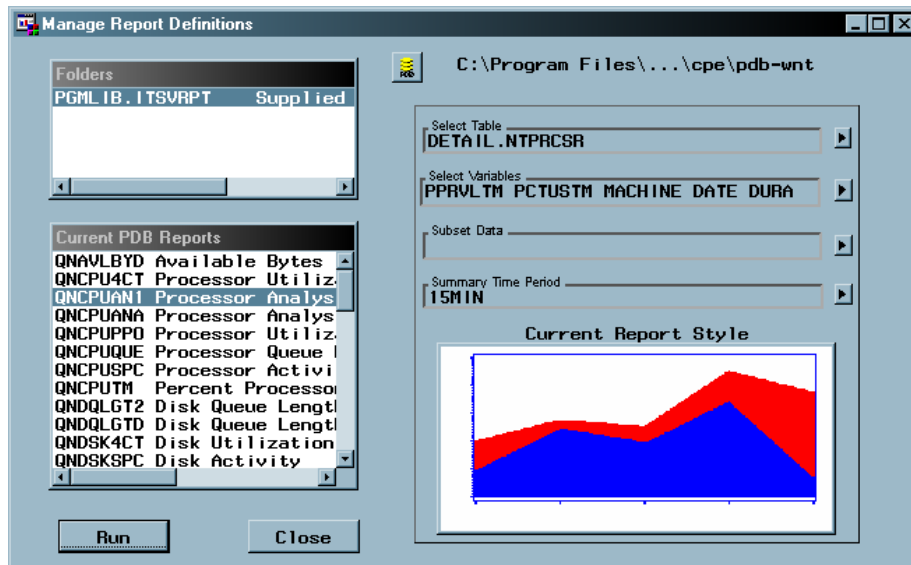□ SAS window

□ External file

□ Web output files.

Typically, if you are using a report tool that enables you to dynamically select attributes and analyze a subset of your data, you can only save or print the report output. This is the case for tools such as Graph Data and Data Visualization. However, the Manage Report Definitions window enables you to create and save the report definition that is used to create the report output. When you specify output options for a report definition, those options are saved with the definition. You can use the report definition to re-generate the same report or an updated version of the report at a later time.

For specific output formats for each report tool, refer to the information about that tool later in this chapter.

A Closer Look at the Report Tools

The Manage Report Definitions Window

The **Manage Report Definitions** window enables you to create and work with reports (and report definitions) that can be run locally or remotely, and either interactively or in batch. This is the only reporting interface within SAS IT Resource Management that enables you to create report definitions that you can save and rerun at a later time.

Through the Manage Report Definitions window, you can accomplish these tasks:

□ Access supplied report definitions for many data sources, and use or modify those report definitions to analyze your own data

□ Create, run, and save custom report definitions

□ Create report definitions that will produce reports that you can display on your local Web site

□ Generate reports from supplied and custom report definitions

□ Create and remove report folders

□ Manage the location of stored report definitions so that they can be accessed by others at your site

□ Create and save report definitions (macro calls) that you can run in batch.

Remember that within the Manage Report Definitions window, there are two parts to the actual report: the report definition, and the report.

The *report definition* defines report attributes such as the data to be used, the PDB, the report style and variables, where to send the output (the actual report), and much more. The *report* is what you actually see when you run that definition.

You can create new report definitions in several ways:

□ Modify a supplied report definition and run it with a demonstration PDB or with your data. You can also save the modified definition as a new custom report definition.

□ Create an entirely new report definition by using the Manage Report Definitions window.

□ Create the report definition by using the report macros.

Note: For more information about running report definitions in batch, refer to the SAS IT Resource Management macro reference documentation. △

Creating and Running a Report Definition

To open the Manage Report Definitions window, start SAS IT Resource Management, and select the **Reporting** tab from the main application window. When the **Reporting** tab displays, select **Manage Report Definitions**.

From this window, you can create a report definition that identifies these report attributes:

□ the folder in which to store the report

□ the report name

□ the table and the data in that table (the level) to use for the report

□ the variables on which to report

□ a subsetting criteria for the report (if you want to report on a subset of your data)

□ a time period on which to report

□ the report style, such as plot, chart, spectrum plot, or other

□ options specific to the report type that you have selected

□ the output format or location for the report, such as a .GIF file, a SAS catalog, or the Graph window

□ and much more.

When you have identified these items for your report definition, you can select **Run** in the Manage Report Definition window. This will run your definition and create your report.
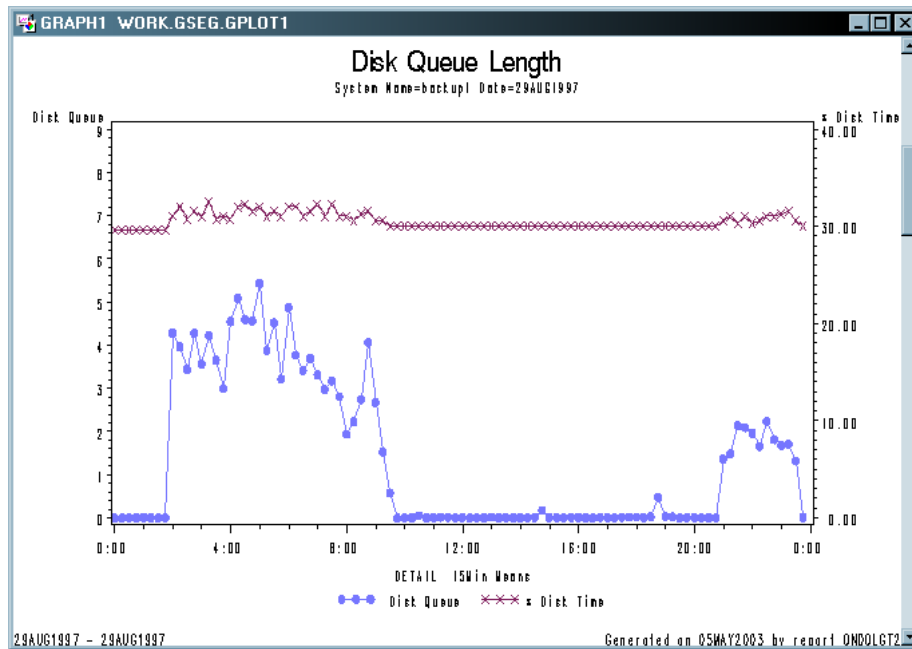
The Manage Report Definitions window also enables you to save the report definition so that you can rerun it and create a report in the future, either within the window interface or in batch. You can also save the report definition as SAS statements and a macro call, as shown in the following sample preview window, so that you can run the report definition in batch mode. Report definitions are especially useful for production reports that you want to rerun each time you process new data.

After you run the report definition, you can save the report as an image in a SAS catalog. You can also export or save the report in a variety of graphic formats (BMP, GIF, TIFF, PS, JPG, and more).

You can perform a variety of tasks from the Manage Report Definitions window. For additional details on completing this window, refer to the topic *Using the Manage Report Definitions Tool* in the chapter *Report Concepts and Tools* in the SAS IT Resource Management User's Guide.

Graph Data

Note: SAS IT Resource Management's Graph Data tool utilizes the SAS Graph-N-Go application. You can access it from the SAS IT Resource Management GUI for UNIX and Windows. △

Graph-N-Go is an interactive facility for graphing and for analyzing your data. Access this Graph-N-Go application by pressing **Graph Data** at the bottom of the Reporting tab. The Select Table for Graph-N-Go Application window opens.
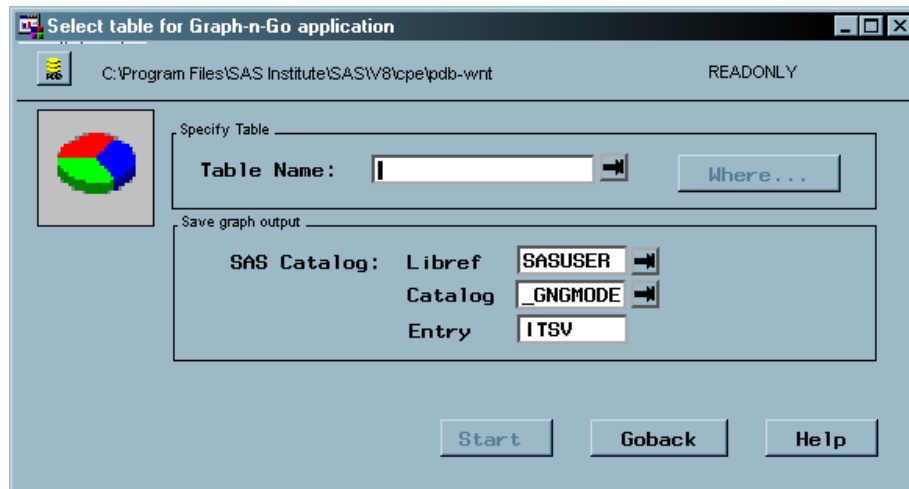
Specify the table that you want to work with. If you want to further subset your input, you can specify a where clause. Click on the **Where** button to enter a subsetting filter.

Note: The **Start** button is ungrayed after you enter the name of the table. △

You can save your work (both the data set models and the report) for future editing. Specify where you want to save your work in the **Save graph output** boxes. By default, the report will be saved in SASUSER._GNGMODE.ITSV.GFORM). You can also export the report to other file types such as HTML. To do this, select the File pull-down menu within the Graph-N-Go application. Then, select the export function.

Press **START** to view the Graph-N-Go window.

The Graph-N-Go window consists of two panes with icons along the left side of the window. The upper pane lets you specify the data that you want to model. It initially contains an icon for the data that you selected in the earlier window.
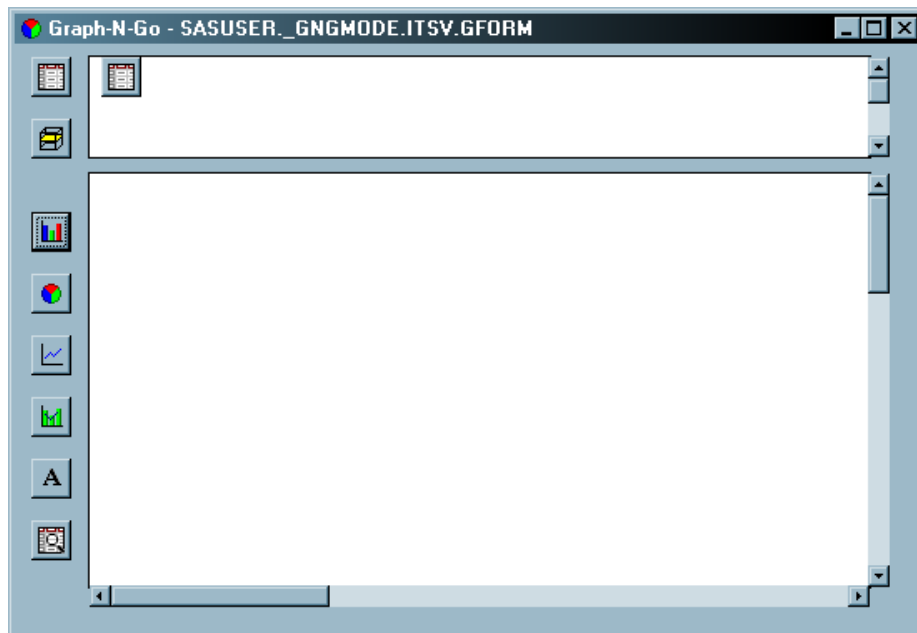
You can add additional data in two ways:

☐ **From a SAS Data Set:** Select the upper icon next to the top pane. The SAS Data Set Selection window opens.
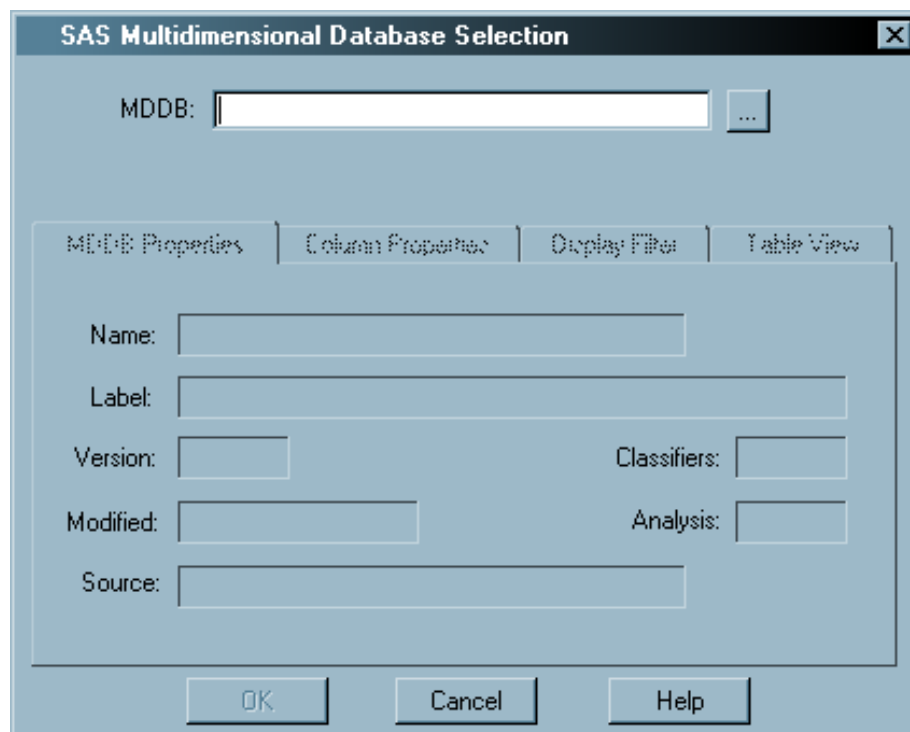

Click the icon to the right of the **SAS Data Set** entry field. The Select A Member window opens. In the **Member_Type** field, select **SQL/Data △ Step/DBI**

Views (VIEW), if it is not already selected. Then, double-click to select the data's library name and view name.
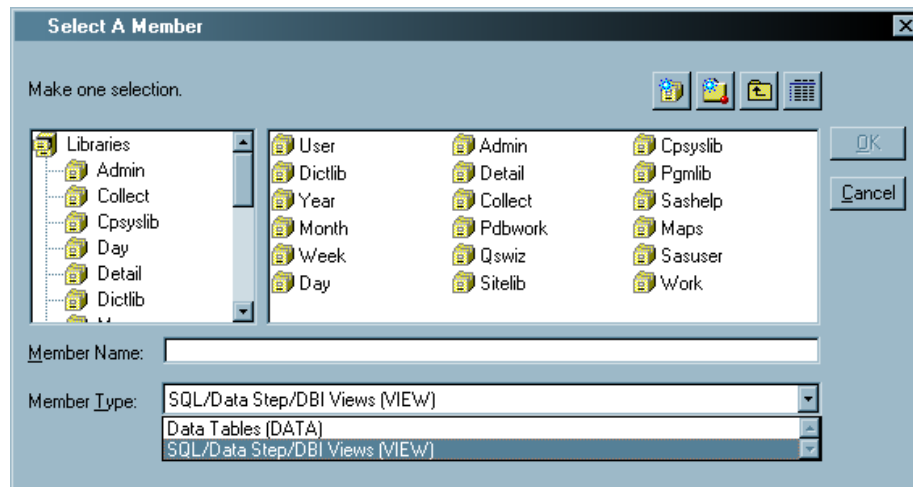
The data model icon for the data that you selected will appear in the upper pane. Double-click on this icon to view the properties of your data or to subset the data further.

☐ **From a Multidimensional Database:** Select the lower icon that is next to the top pane. This lets you select data from a multidimensional database (MDDB).

Click the icon to the right of the **MDDB** field. The Select A Member window opens. Select the library and then double-click to select the MDDB that you want to work with. The data model icon for the data that you selected will appear in the upper pane. Double-click on this icon to view the properties of your data or to subset the data further.

After you have specified all the data you want to work with in the upper pane, you can specify the graphical viewers you want to use to display the data. The icons next to the lower pane represent the different graphical viewers that you can use. To select a viewer, click its icon. (You can select bar charts, pie charts, and so on.) Drag the viewer to an empty part of the lower pane and click. Next, drag and drop a model from the upper pane onto the viewer. The data is displayed according to the default properties for that viewer. You can modify the display by putting your mouse pointer over the graph, clicking your mouse selection button, and choosing from the selection menu. Multiple graphs and/or tables can be displayed on the lower pane.

To facilitate future editing, you can save any report that you develop in a SAS library. (By default, the report will be saved in SASUSER._GNGMODE.ITSV.GFORM.) You can also export the report to other files types such as HTML. (See the export function in the **File** pull-down menu within Graph-N-Go.)

For more information about this topic, refer to the online help within the Graph-N-Go application.

Exception Reports

The Perform Exception Analysis tool helps you create exception rules that you can evaluate (run) to discover exceptions in your data. For example, you might want to identify conditions when disk usage was exceptionally high, or when queue length or wait times were exceptionally long.

You can create folders in which to store the exception rules, and you can also evaluate a single rule or all rules in a folder. After you evaluate rules in a folder, only the rules that apply to tables in the current PDB are evaluated. Rules that do not apply to tables in the current PDB are ignored. When you evaluate the exception rules, you can view reports on the exceptional data.

To create and evaluate exception rules, follow these steps:

1 Select **Perform Exception Analysis** on the **Reporting** tab in the SAS IT Resource Management client.

The Exception Reporting window opens. This window contains a list of the folders that contain exception rules. The *PGMLIB.CPXRULE* folder contains exception rules that are supplied with SAS IT Resource Management. You can create a new folder, add folders to this list, create a new rule, or modify an existing rule. To create and save a new exception rule, you must have write access to the folder and the location where you want to store the rule.
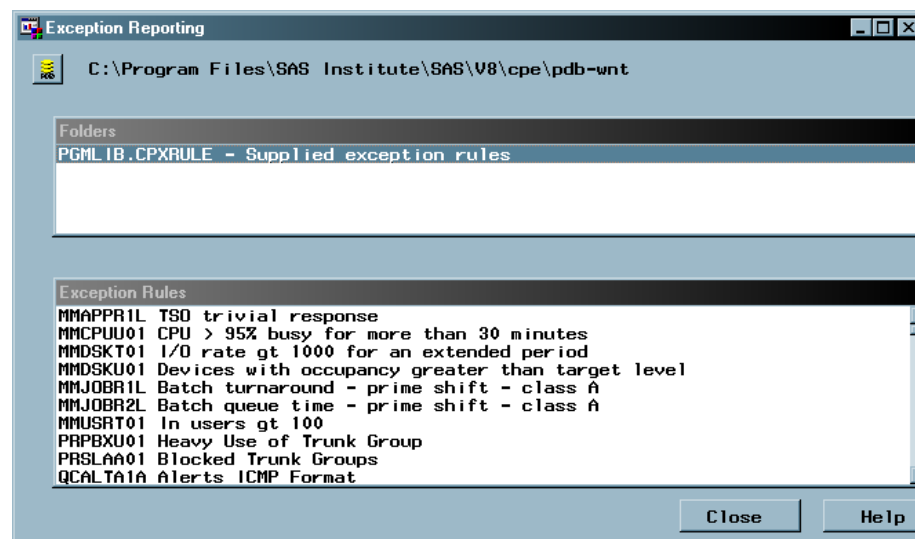
2 Before you create or add folders or rules, verify that the active PDB is the one that contains the data with which you want to work. The PDB name is located in the top portion of the Exception Reporting window.

3 From the Exception Reporting window, you can create a new exception folder by selecting the following menu path: **File ▶ Folders ▶ Create/Add**.

Type or select the library name, the folder name, and a description in the Create or Add Folder window. Select **OK** to create the folder and to close the window.

4 To create a new rule, select the folder in which to store the rule. Then, select the following menu path: **File ▶ Rules ▶ New**.

You can also delete rules and folders by selecting the appropriate items from the **File** menu.

5 Complete the Properties for Exception Rule window to identify the rule name, description, location of the data against which the rule will be evaluated, the rule expression, subsetting criteria, and more. For assistance in completing this window, select **Help** in the window.
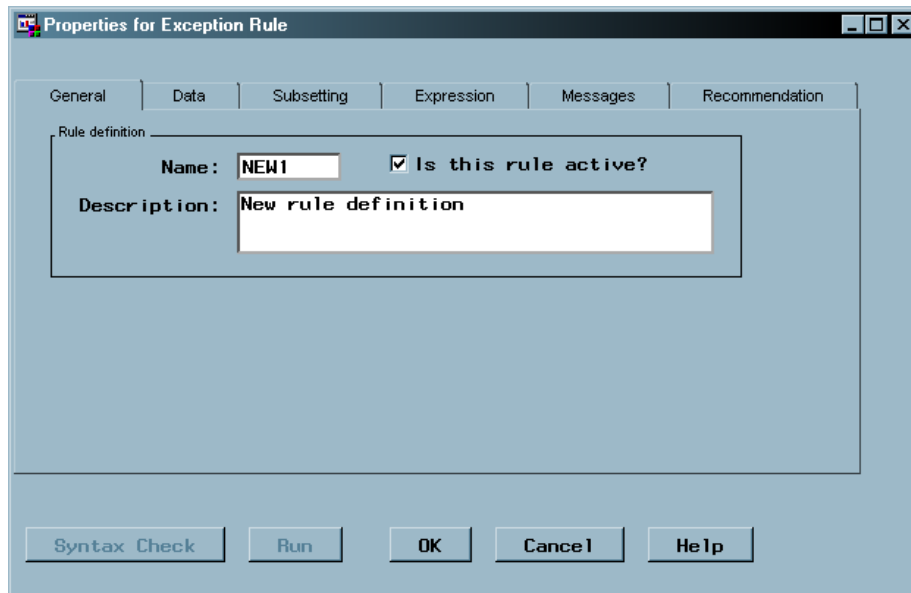
6 To evaluate all the rules in an exception folder, select the folder from the list in the Exception Reporting window. Then, select: **File ▶ Folders ▶ Evaluate**.

7 To evaluate a single rule in a folder, highlight the folder that contains the rule, and select the rule from the list. Then, select **File ▶ Evaluate**.

8 If any rules are found to be true, based on the rule expression, you can view the results of the exception evaluation by selecting: **Locals ▶ Results**.
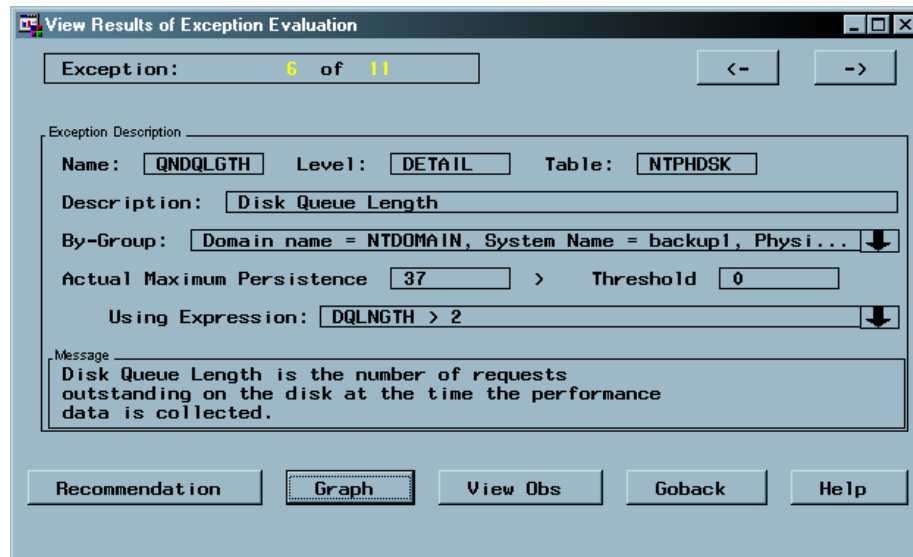
The View Results of Exception Evaluation window provides information about the exceptions in each By group. For example, if you evaluate a rules folder and seven exceptions are found, this means that there are seven By groups with exceptions.

9 Select **Graph** in the View Results of Exception Evaluation window to view the exception in the graphic context of other data. For example, the following graph displays each data point in the By group (Domain, System, and Instance).
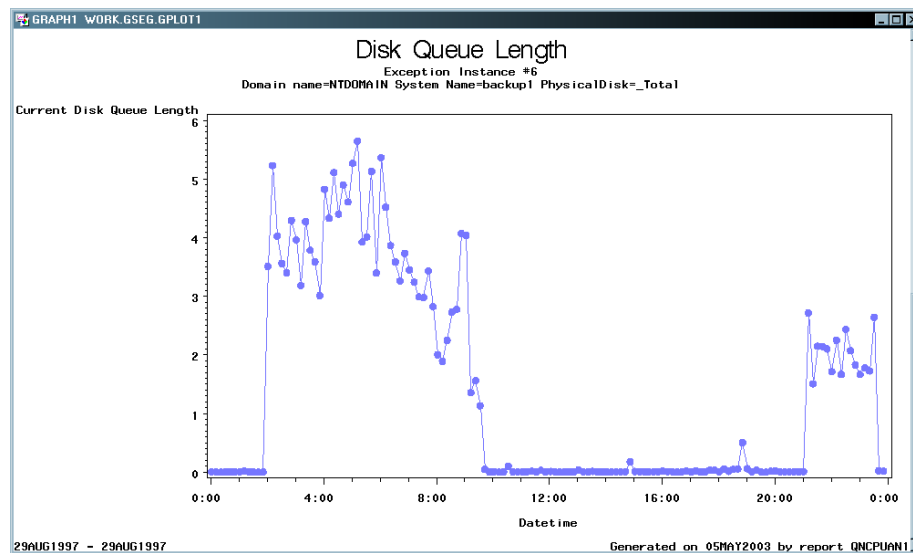
You can also select **View Obs** to view the text context for the exceptions.

10 To save the exception report image as a .BMP file, select the following path within the interface: **File ▶ Export as Image**.

Exception rules must be defined within the GUI. However, they can then be evaluated interactively (by using the GUI), or in batch mode (by using the %CPEXCEPT macro). In batch mode, you can also use the %CPXHTML macro to create exception reports that you can display by using a Web browser.

Data Visualization

The Data Visualization tool enables you to explore and analyze data through interactive 3-D plots, rotating plots, histograms, scatter plots, and charts. By using the point-and-click interface, you can quickly and easily examine relationships between performance metrics, perform trend analysis, and visualize your performance data in one or more dimensions. You can analyze distributions, calculate correlations and principal components, and fit models within the framework of the generalized linear model.

This tool is integrated with SAS IT Resource Management's dynamic graphics and advanced data analysis tools. The tool provides a valuable addition to your performance analysis needs.

To create a graph by using the Data Visualization tool, follow these steps:

1 Select **Data Visualization** on the **Reporting** tab in the main SAS IT Resource Management window.

2 The Start Data Visualization Session window opens.
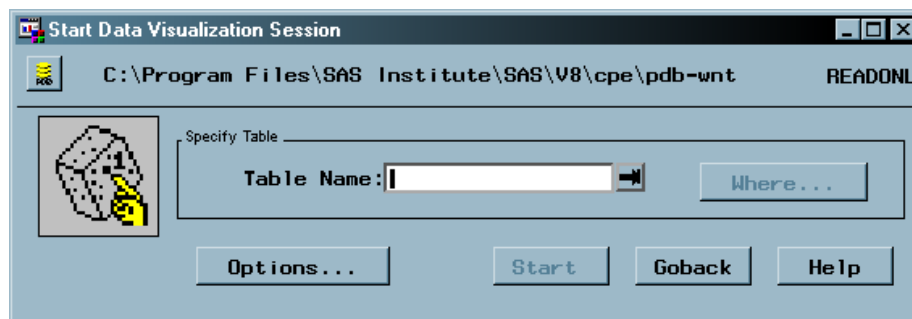
From this window, you can select the table that you want to analyze.

Note: To browse the list of tables, click the arrow that is next to the **Table Name** field. △

Select the table and click **OK**.
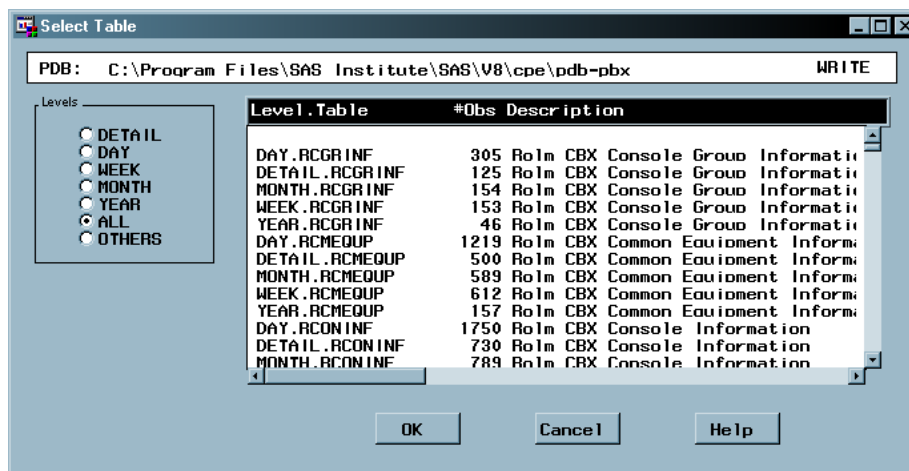
3 The Start Data Visualization Session window is displayed again, showing the table that contains the data that you selected. The number of observations and the number of variables display in the top left corner of the window.

From the **Analyze** menu, select the report style, and then select the variables to use in your report. When the report displays, you can dynamically add or modify other report attributes such as the color, tick marks, legends, and more, or create additional reports, as shown in the following sample report.
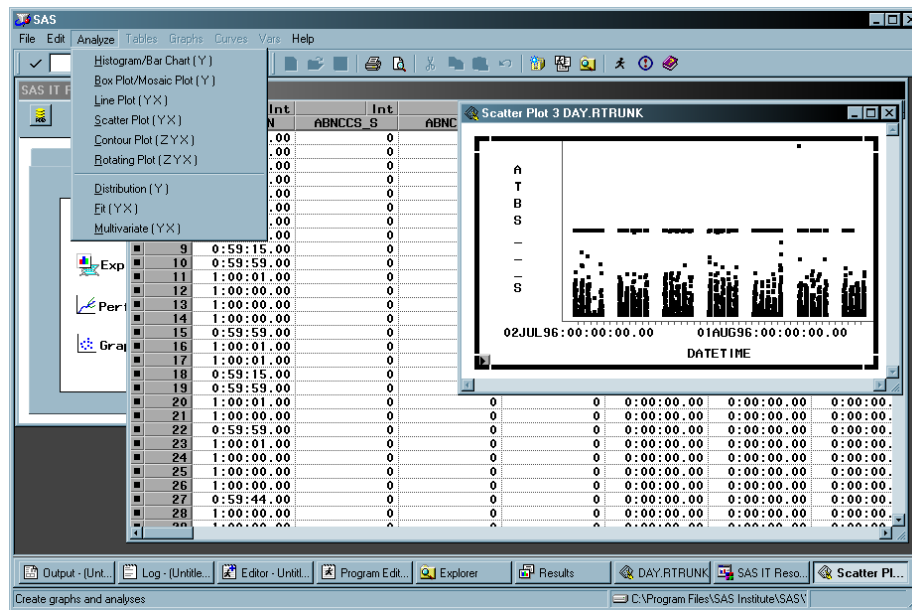
4 After you create the report, you can print the report directly from the Data Visualization window or you can save the report in one of the following formats:

   □ Save the report output to a graphics catalog — Graph or Image

   □ Save the report output to a graphics file — BMP, GIF, PBM, PS, or TIFF formats.

From the Data Visualization window, you cannot save the specifications to re-create this report. However, you can save the steps used to create the report and then recall and rerun those steps in the future.

This reporting tool is well suited for ad hoc analysis through interactive plots and charts because it helps you to create reports on your production data dynamically.

Displaying Reports on the Web

SAS IT Resource Management enables you to create a report definition and specify an output mode of *Web*. This enables you to view the reports by using a Web browser, such as Microsoft Internet Explorer or Netscape Navigator. You can select the *Web* output mode through the window interface (on UNIX or Windows) or you can specify the *Web* output mode on your calls to report macros (on all platforms).

When you run the report definition, your reports are saved as images. This output format also generates text output files that utilize a markup language called HTML. The HTML files reference your report images, and when you display the HTML file by using your Web browser, your report images are also displayed.

When you select an output mode of *Web*, you must also set options for the Web reports, such as where your HTML output and images are stored, the size of your images, the style of your report, and more. You can set these options with the macros that run in batch, for example, the %CPRUNRPT macro. For more information about the report macros, see "Creating Web-Based Reports by Using Report Macros" on page 63.

Creating Web-Based Reports by Using the Window Interface

To create a new Web report from the window interface, follow these steps:

1 Create a report definition for your report (as you would for any type of report) by completing the Manage Report Definitions window.

2 After you have created the report definition in that window, follow this path from the Manage Report Definitions window: **Locals ▶ Report Output Options**.

This opens the Report Output Options window.

3 Select `Web` from the Report Output Options window and then select `Set Attributes` to open the Generate Web Output window.
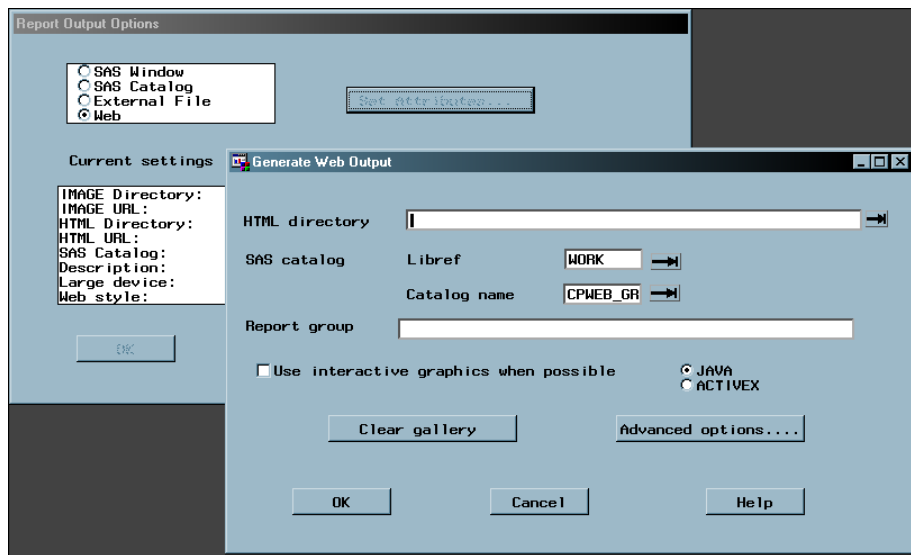
4 Complete the Generate Web Output window in order to define additional output options for your Web report. The options in this window enable you to specify additional report options, such as where your report images are stored, where your HTML files are stored, the Web report style, and more.

5 Select `Run` in the Manage Report Definitions window. This runs the report definition and displays the resulting HTML files by using your Web browser, as shown in the following sample report output.
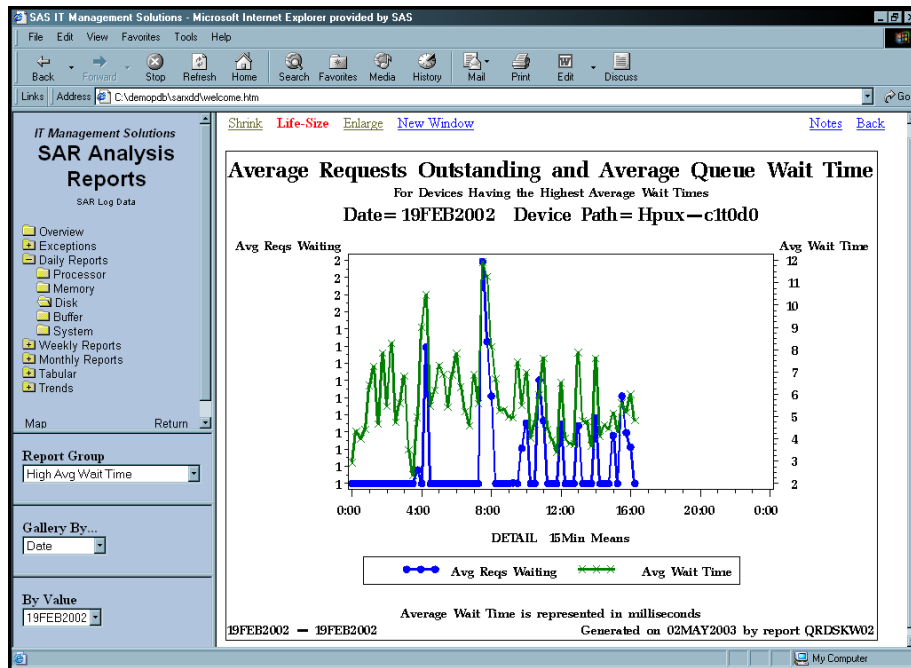
Creating Web-Based Reports by Using Report Macros

By using the batch report macros, you can specify the OUTMODE= parameter with a value of *WEB* in order to create reports that you can display by using a Web browser. For more information about report macros, refer to the SAS IT Resource Management macro documentation.

To create a Web report by using batch macros, follow these steps:

1 Determine which style of report you want to create. Then, select the appropriate report macro to create that report.

2 Create a report definition by coding the macro in your program and then by submitting or running the program in batch mode. When you code the macro, you must specify OUTMODE=WEB as well as the additional Web parameters available on the macro. These parameters set other Web report options, such as where your report images are stored, where your HTML files are stored, and more.

3 If you create the report on a z/OS server by specifying a UNIX File System area, the report will be written directly to that area. If you create the report on a z/OS server without specifying a UNIX File System area, the report will be written to a PDS. You can then FTP the report from the PDS to UNIX or Windows to view the report with a Web browser.

C H A P T E R

5
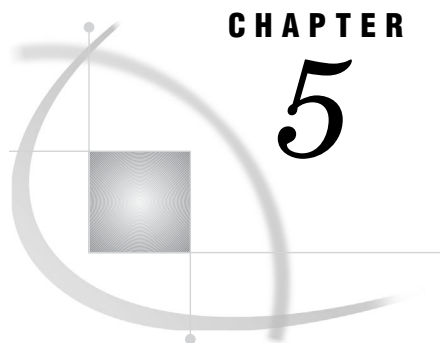
Additional Documentation and Resources

Overview

SAS IT Resource Management provides many other resources to assist you in setting up and using this product, and managing your IT resources. These resources include:

- □ online documentation — window help and SAS System help
- □ Web documentation — what's new, the User's Guide, the Macro Reference document, a glossary, and the Server Setup Guide.
- □ printed documentation — course notes for the Interactive Reporting course and the Administration and Batch Reporting course.
- □ training and consulting
- □ technical support
- □ e-mail forum
- □ chargeback functionality through SAS IT Charge Management.

Each of these resources is described further in the following sections of this chapter.

Web and Online Documentation

A variety of Web information about SAS IT Resource Management is available from the http://support.sas.com/documentation/onlinedoc/itsv Web page:

- □ What's New in SAS IT Resource Management — in this release and several previous releases
- □ Getting Started with SAS IT Resource Management— a copy of the document that you are reading now
- □ SAS IT Resource Management User's Guide — how-to documentation for interactive mode, with pointers to the corresponding work in batch mode

  ☐ General

  ☐ Setup

  ☐ Administration

  ☐ Reporting

 ☐ SAS IT Resource Management Server Setup Guide — a collector-specific supplement to the setup cases in the User's Guide

 ☐ SAS IT Resource Management Macro Reference document — working in batch mode

 ☐ SAS IT Resource Management Glossary

To access this Web page, use a frame-capable version of a Web browser such as Microsoft Internet Explorer or Netscape Navigator.

 Within the SAS IT Resource Management GUI for UNIX and Windows, online documentation includes window help and also an HTML page named *Other ITRM Documentation*, which has access to the following documentation (in addition to copies of some of the other documentation on the Web):

 ☐ Collector Updates for this Release - contains details on the updates that have been made to the supplied collector support at this release

 ☐ SAS IT Resource Management on the Web - includes links to Online Books, White Papers, Brochures and Documentation

 ☐ Server Setup Guide - describes how to set up your SAS IT Resource Management server for a variety of data collectors.

 ☐ QuickStart Examples — contains sample report output from the QuickStart Wizard

 ☐ Migration - explains how to migrate your SAS IT Resource Management environment from one SAS version to another.

To access Other ITRM Documentation, from the main window of the SAS IT Resource Management GUI for UNIX and Windows or the SAS IT Resource Management client GUI, select **OnlineHelp ▶ Other ITRM Documentation**.

 Within the SAS IT Resource Management GUI for z/OS, online documentation includes window help and also a Help menu, which has access to the following documentation (in addition to copies of some of the documentation on the Web):

 ☐ What's New in SAS IT Resource Management

 ☐ The Help Index

 ☐ Message Help

 ☐ MXG Online Documentation.

To access the Help menu, from the main menu of SAS IT Resource Management, select `Help`.

Setup

 The main instructions for setting up your data source and SAS IT Resource Management so that data flows from the data source into the warehouse and reports is in the Setup part of the User's Guide. In the User's Guide, follow the setup case that applies to data from your collector, and use the Server Setup Guide to supplement the setup case with the collector-specific information for your collector's data.

 The Setup part of the User's Guide also contains documents about the MIB to Dictionary Compiler, MXG-based Tables, the Generic Collector Facility, and Migration from SAS 6 to SAS 8.

 Extensions to SAS IT Resource Management, such as use of process exits, alternative methods of duplicate data handling, and creation and installation of collector packages, are covered in the Administration part of the User's Guide.

SAS System Documentation

The SAS System includes many software products and solutions in addition to SAS IT Resource Management. SAS System documentation is shipped online with your SAS software. You can access SAS documentation from the SAS IT Resource Management client GUI by selecting **OnlineHelp** ▶ **SAS IT Resource Management Help**. Scroll in the left frame until you find the category of help that you want to see, and select it. The help displays in the right frame.

You can also order hardcopy SAS books from SAS Book Sales (1-800-727-3228).

Printed Documentation

You can print the documents that are in PDF format on the http://support.sas.com/documentation/onlinedoc/itsv/ Web page.

Course notes for SAS IT Resource Management courses (SAS IT Resource Management: Interactive Reporting, and SAS IT Resource Management: Administration and Batch Reporting) are available from the SAS Publications Catalog. On the http://support.sas.com Web page. select **Bookstore** ▶ **Complete Catalog** ▶ **Course Notes**.

You can also order additional SAS books in print from the complete catalog. The following SAS documentation can help you get started with SAS or to help you learn to use SAS and SAS IT Resource Management on a new platform:

☐ SAS Companion documentation, available for each host (PC, UNIX, z/OS) under the section for the Base SAS product

☐ SAS/GRAPH product section — Introduction and Reference books

☐ SAS/AF product section — Usage and Reference books

You can order these and other books on the Web, or by using the Documentation Request form that is shipped with SAS IT Resource Management, or by contacting SAS Book Sales (1-800-727-3228).

Training

The following SAS IT Resource Management courses are available:

☐ Interactive Reporting — discusses SAS IT Resource Management basic concepts, how to use various reporting tools, and how to create and manage your reports through the Manage Report Definitions window and other online reporting tools.

☐ Administration and Batch Reporting — provides information for the SAS IT Resource Management administrator about creating and managing your PDB, site options, and more. The information includes interactive administration, batch administration, and batch reporting.

For more information about training, call 1-800-333-7660.

Technical Support

Technical Support is available to all sites licensing software from SAS. The document "Overview of Technical Support Services" is provided with your software shipment.

If you have questions or problems about your licensed software, you might want to explore the SAS Institute Technical Support Web site at

```
http://support.sas.com/techsup/intro.html
```

The Technical Support Web site offers SAS Notes, Technical Support documents, FAQs, SAS sample programs, phone numbers, e-mail support, and more that might answer your questions. This Web site also provides an online mechanism for reporting problems. SAS Institute also maintains several other electronic services for tracking problems and questions for Technical Support. These electronic services are available 24 hours a day.

If you want the question or problem handled by phone, please have your site's SAS Installation Representative or SAS Support Consultant call the SAS Technical Support Division. Sites in the U.S. and Canada can call (919) 677-8008. Other sites should check the Technical Support Web site or contact their SAS Installation Representative or SAS Support Consultant for the nearest SAS Institute office.

SAS IT Resource Management E-Mail Forum

You can exchange information with other SAS IT Resource Management users by subscribing to the ITMS-L listserv. This is an e-mail forum that is available for users to exchange tips, techniques, and information about SAS IT Resource Management. To subscribe, send an e-mail message to *listserv@vm.sas.com*, leave the subject line blank, and type

```
subscribe itms-l your-first-name your-last-name
```

in the body of the message.

To send mail to the list, send e-mail to **itms-l@vm.sas.com**.

SAS IT Charge Management Product Information

SAS IT Charge Management ships as a separate product, but it is designed for use with SAS IT Resource Management software to provide chargeback applications for service data. SAS IT Charge Management evaluates and recovers service costs through accounting and budgeting methods and site-specific standards. If your site is interested in invoicing for IT services, ask your sales representative about SAS IT Charge Management.
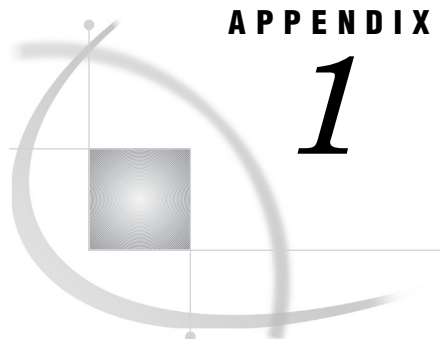
Recommended Reading

Recommended Reading

Here is the recommended reading list for this title:

- □ *SAS IT Resource Management: User's Guide*
- □ *SAS IT Resource Management: Administrative and Batch Reporting Course Notes*
- □ *SAS IT Resource Management: Interactive Reporting Course Notes*

For a complete list of SAS publications, see the current *SAS Publishing Catalog*. To order the most current publications or to receive a free copy of the catalog, contact a SAS representative at

SAS Publishing Sales
SAS Campus Drive
Cary, NC 27513
Telephone: (800) 727-3228*
Fax: (919) 677-8166
E-mail: **sasbook@sas.com**
Web address: **support.sas.com/publishing**
* For other SAS Institute business, call (919) 677-8000.

Customers outside the United States should contact their local SAS office.

Glossary

activate
to specify that a particular performance data warehouse, site library, or table is the one that you want to use now.

active PDB
the performance data warehouse that you are currently using.

active site library
the library of site-related information that you are currently using.

active table
the table to which the control statements in a stream of %CPDDUTL control statements currently refer.

administration tasks
tasks that manage the SAS IT Resource Management server. Administration tasks typically require write access to a performance data warehouse (PDB), to a site library, or to a program library. See also reporting tasks.

age limit
in a table, a unit of time that describes how long existing data is to be kept in a performance data warehouse (PDB) when new data is processed and reduced. An age limit specification pertains to the data that is already in a particular level of a particular table, and does not pertain to any data that is being read in. In a report gallery, the age limit specification describes how long the reports are to be kept.

aged data
data that is deleted from a level of a table in the performance data warehouse (PDB) by the process task or the reduce task because the data exceeds the age limit for keeping data in that level of that table.

aggregate
another term for summarize.

analysis variable
a numeric variable that is used to calculate statistics or to display values. Usually an analysis variable contains quantitative or continuous values, but this is not required.

API (application programming interface)
a set of software functions that facilitate communication between applications and other kinds of programs, services, or devices.

archive

a set of archive libraries for a specific performance data warehouse (PDB).

archive library

a SAS library that contains the data that was incoming to the detail level of a performance data warehouse (PDB) during a particular execution of the process task. An archive library is one of the libraries in the archive. See also detail-level data.

archive location

for a performance data warehouse (PDB), specifies where the archive is stored, if at least one table specifies archiving. The archive location (root) on z/OS is a partitioned data set (PDS). The archive location (root) on UNIX and Windows is a directory.

archive status

for a table, a value that specifies whether data that was incoming to the detail level of the table is to be copied to an archive.

archived data

data that was incoming to the detail level and which was then copied to an archive library during an execution of the process task.

autoexec file

a file that contains SAS statements that are executed automatically when SAS is invoked. The autoexec file can be used to specify SAS system options, as well as to specify librefs and filerefs for data sources that are used frequently.

availability

the percentage of time that a system is available to respond to the requests of its users.

axis

a one-dimensional line that represents the scale that is used for plotting the values of x, y, or z coordinates.

backload

to process and reduce data that was logged before the most recently logged data in the performance data warehouse (PDB).

backup copy

a complete copy of one or more libraries.

batch job

a unit of work that is submitted to an operating system for batch processing. On z/OS, a batch job is a set of JCL statements; on Windows, a batch job is a task; and on UNIX, a batch job is a background process.

batch mode

a method of executing SAS programs in which a file that contains SAS statements plus any necessary operating environment commands is submitted to the computer's batch queue. After you submit the program, control returns to your terminal or workstation, where you can perform other tasks. Batch mode is sometimes referred to as running in the background. The program output can be written to files or printed on an output device.

browser

See Web browser.

BY variables list

1) in the detail level of a table definition, a list of variables that specifies the sort order and grouping of the observations in the detail level of that table. 2) in a report definition, a list of variables that specifies the sort order and grouping of the data in

the report. 3) in a rule definition, a list of variables that specifies the sort order and grouping of the data in the results. See also exception, CLASS variables list.

calculation status
1) for a formula variable at any level, an indicator that specifies whether the value of the formula variable is to be calculated at that level. 2) for a summary statistic at a summary level, an indicator that specifies whether the value of that summary statistic is to be calculated at that summary level.

capacity planning
the process of measuring performance, availability, and workload volume in a particular computing environment and comparing these measurements over time to baseline values. The object of capacity planning is to predict when resources will need to be upgraded.

catalog entry
See entry type, SAS catalog.

chart
a graph in which graphical elements such as bars or pie slices show the magnitude of data values. The graphical elements can represent one data value or a range of data values.

CLASS variable
a variable that is used to group, or classify, data. The data type of CLASS variables can be either character or numeric. Class variables can have continuous values, but they typically have a few discrete values that define the classifications of the variable.

CLASS variables list
1) in the day, week, month, or year level of a table definition, a list of variables that specifies the sort order and grouping of the data that is summarized into observations in that level of that table. There is one CLASS variables list for each summary level of each table. 2) in a report definition, zero, one, or more variables that may interact with other variables (such as the Group, SubGroup, or Stack variables) to specify the grouping of the data in the report.

client
See SAS IT Resource Management client.

collector
another term for data collector.

collector name
one of the two identifiers that specifies a particular set of collector-support software. The other identifier is tool name. For example, the pair of identifiers specifies which collector- specific staging code to run during the process task.

collector support
a set of software components that provide a variety of functions that enable SAS IT Resource Management to work with data from a particular collector. Each collector-support component is identified by both a collector name and a tool name.

combined view
a view that provides access to variables from two or more other views.

continuous values
values of a variable that theoretically could indicate an uncountable number of possible values. For example, a variable that measures the temperature of water in degrees Celsius under normal conditions is continuous because it can have any value from 0 to 100, even though the thermometer that is used is not able to measure temperatures beyond a particular level of decimal precision.

control data set
one of the data sets that the duplicate-data-checking macros maintain and use to determine whether data from a log has already been processed.

control statement
an instruction or a string of keywords that you provide to the data dictionary utility in order to direct the utility's actions in batch mode. For example, control statements can direct the utility to copy a table definition from the supplied table definitions in the master data dictionary to the active performance data warehouse (PDB), to create and print table definitions, and to perform a variety of other functions.

cron
a UNIX utility for executing batch jobs on a periodic basis.

current site library
another term for active site library.

current table
another term for active table.

daemon
a process that starts and waits either for a request to perform work or for an occurrence of a particular event. After the daemon receives the request or detects the occurrence, it performs the appropriate action. If nothing else is in its queue, the daemon then returns to its wait state.

data collector
software that collects logged data that is related to IT services. The data collector could be a facility of the operating system, a product that is provided by a third-party vendor, or software that is written by consultants or by SAS IT Resource Management customers.

data dictionary
See master data dictionary, PDB's data dictionary.

data dictionary utility
a macro (%CPDDUTL) that is used for adding, changing, deleting, or checking the information in the data dictionary of the performance data warehouse (PDB). The utility can also generate a preliminary version of a table definition (and its associated variable definitions) from a SAS data set or view, from a character-delimited external file, or from other types of files.

data reduction
the summarization of data for a day, week, month, or year. The observations for the current week, month, and year contain week-to-date, month-to-date, and year-to-date data. Data reduction decreases the amount of memory that is required for storing data, but it still preserves information that can be used for analyzing trends.

data source software
software that logs data that is related to IT services.

data type
an attribute of a variable. In SAS IT Resource Management, a variable can have a data type of either character or numeric. See also interpretation type.

data visualization
the act of representing data visually or graphically. Charts, plots, and graphs are some results of data visualization.

date format
in SAS software, the instructions that tell SAS how to write numeric values as date values.

datetime format
in SAS software, the instructions that tell SAS how to write numeric values as datetime values.

day-level data
performance data that is summarized for an entire day. Day-level data is summarized from detail-level data.

derived variable
a variable whose values are calculated from regular variables, constants, or both. See also regular variable, formula variable.

detail-level data
in general, performance data that has not been summarized. (For data collectors such as SAP that log an extremely large volume of data, detail-level data can be partially summarized.) Some of the data might have been changed by the staging code, by user-written exit routines for the process task, by partial summarization, or by some other means. See also day-level data, week-level data, month-level data, year-level data.

duplicate data
1) observations that have identical values in all of the BY or CLASS variables. 2) observations that come from the same log. The meaning that is intended is explained when the term is used.

duplicate-data-checking macros
a set of macros (%CPDUPINT, %CPDUPDSN, %CPDUPCHK, and %CPDUPUPD) that control whether data from a log can be reprocessed into a performance data warehouse (PDB). See also control data set.

duration
1) in records that represent interval data but which do not contain a value for the length of the interval, a value in seconds that represents the length of time between records. 2) in records that represent interval data and which do contain a value for the length of the interval, that value in seconds. 3) in observations for interval-type tables, the variable that contains the length of the interval in seconds.

entry type
a characteristic of a SAS catalog entry that identifies the catalog entry's attributes and structure to SAS software. When a catalog entry is created, SAS automatically assigns the entry type as part of the name.

ETL (extract, transform, load)
another term for process.

event-type table
a table in which each observation represents an event. The timestamp on the observation represents the date and time of the event. See also interval-type table.

exception
an observation that results from evaluating a rule. Each set of unique values of the rule's BY variables represents a group of observations against which the rule ran. The total number of exceptions indicates the number of groups for which the rule is true, not the number of individual observations for which the rule is true. See also results.

exit code
additional code that is to be executed at a specified point in the supplied software.

exit point
a point in the supplied software where additional code runs (if additional code is provided).

external detail
a view that enables data that is formatted like detail-level data to be reduced directly into the summary levels of the performance data warehouse (PDB) and therefore to bypass the detail level in the PDB. The data to be reduced might or might not be in a PDB.

external file
a file that is created and maintained by a host operating system or by another vendor's software application. Either data or stored SAS statements can be read from an external file. Data can be written as output to an external file. A SAS data set is not an external file.

external name
1) when referring to a SAS IT Resource Management table, the name of the object that contains the data that is to be merged into the table. The object can be a staged data set, a staged view, or a character- delimited file. 2) when referring to a SAS IT Resource Management variable, the name of the corresponding variable in the staged data set, in the staged view, or in the character-delimited file.

fileref (file reference)
a short name (or alias) for the full physical name of an external file. A SAS FILENAME statement maps the fileref to the full physical name.

folder
a repository for objects that are of the same type. For example, report definitions are stored in one or more report definition folders, rule definitions are stored in one or more rule definition folders, and palette definitions are stored in one or more palette definition folders.

formula variable
a variable whose values are calculated from regular variables, derived variables, other formula variables, and/or constants. The values are calculated when they are needed; they are not stored. The formula's source statements are stored. See also regular variable, derived variable.

fully populated report
a report that contains as much data as the report definition allows.

fully populated table
a table that contains as much data as the age limit specifications allow.

gallery
all the reports (and the associated left and right frames) that are visible in a Web browser when a particular welcome.htm file is selected.

generic collector
one of the data collectors for which no table definitions are supplied.

Generic Collector Facility
a component of SAS IT Resource Management that enables users to process data into a table for which the user provides the table definition. The table is typically defined by using the GENERATE SOURCE control statement.

global macro variable
a macro variable that can be referenced in any referencing environment in a SAS program, except when there is a local macro variable that has the same name. A global macro variable exists until the end of the session or program. See also macro variable.

graph report
a report that uses high-resolution graphics. See also text report.

graphics option
in a SAS GOPTIONS statement, an option that controls some attribute of the graphics output. The specified value remains in effect only for the duration of the SAS session. Some graphics options override parameters that have been specified for a graphics output device.

graphics output device
any terminal, printer, or other output device that is capable of displaying or producing graphical output.

graphics stream file (GSF)
a file that contains device-dependent graphics commands from a SAS/GRAPH device driver. This file can be sent to a graphics device or to other software applications.

GSF
See graphics stream file (GSF).

GUI
graphical user interface.

host
an operating environment that is identified by a name. For example, a domain name or an IP address can be used to identify a computer or a router, and an LPAR name can be used to identify an LPAR.

ID variable
a variable that contains an alternate identifier for the data in a CLASS variable. For example, the value of a CLASS variable could be a device address, and the value of the corresponding ID variable could be the name of the device. ID variables are useful for identification because their values rarely change.

ID variables list
a list of ID variables. There is one ID variables list for each summary level of each table.

index variables list
a list of variables for which a simple (non-compound) index is to be built. There is one index variables list for each level of each table. If a variable is indexed, a WHERE expression that includes that variable, such as a WHERE expression that subsets data for a report definition or a rule definition, typically runs more quickly.

install a table
to store a table definition and its associated variable definitions (as well as override control statements, if any) in the master data dictionary, so that the table definition can be used for other performance data warehouses (PDBs) as if it were a supplied table definition.

install collector support
to integrate user-written collector-support software into SAS IT Resource Management so that it looks and works like the collector-support software that is supplied with SAS IT Resource Management.

Internet
 the proper name of the super-network that connects many smaller networks around
 the world and which enables all the computers on those networks to exchange
 information. All of the computers on the Internet use the same set of communication
 protocols, called TCP/IP. The World Wide Web uses the Internet as an electronic
 highway. See also internet, intranet.

internet
 any group of interconnected networks that use Internet Protocol (IP) for exchanging
 information. The individual networks are not necessarily owned by the same
 company or organization. The Internet is the biggest example of an internet. See
 also intranet.

interpolate
 to estimate values that are between two or more known values.

interpretation type
 an attribute of a variable that further defines the data type. For example, the data
 type Numeric has interpretation types such as Count, Time, Datetime, and Gauge.
 The interpretation type determines the default settings of other attributes of the
 variable, such as maximum length and format.

interval-type table
 a table in which each observation represents an interval of time. The timestamp in
 the observation typically represents the beginning of the interval, and the duration
 in the observation represents the length of the interval. See also event-type table.

intranet
 a private network of interconnected networks within a single company or
 organization. These networks use the same TCP/IP communication protocols that are
 used on the public, worldwide Internet, but they are either closed to the Internet or
 they are connected to the Internet via a firewall. A firewall enables internal users to
 access the Internet, but prevents unauthorized outsiders from accessing the intranet.
 In other words, intranets enable companies or organizations to share information
 without making that information available to everyone who has access to the
 Internet. See also Internet, internet.

IT
 information technology.

IT services
 services that an IT organization provides to customers. IT services could include
 support for computer systems, networks, applications, telephones, fax machines, and
 Web pages, as well as other types of software and hardware services.

keep status
 1) for a table, a value that indicates whether the table is to be used or ignored. 2) for
 a variable in a table, a value that indicates whether the variable is to be used or
 ignored.

keyword parameter
 a parameter that is identified by name. For example, in the macro call
 %CPDUPDSN (SOURCE=WWW);, 'SOURCE' is a keyword parameter that specifies
 the input for the macro %CPDUPDSN. Keyword parameters can be used in any
 order on the macro call, but they must follow any positional parameters.

LAN
 local area network.

level
a logical grouping of all data that is consolidated for a specific unit of time. Within a table, there are five levels in which data is stored: detail, day, week, month, and year.

libref (library reference)
a short name (or alias) for the full physical name of a SAS library. A SAS LIBNAME statement maps the libref to the full physical name. A libref is the first part of a multi-level SAS filename and indicates the SAS data library in which a SAS file is stored. For example, in the name SASUSER.ACCTS, SASUSER is the libref, and ACCTS is a file in the library that the SASUSER libref refers to. See also SAS library.

local PDB
a performance data warehouse (PDB) that is located on a disk that is attached to the local computer or which is on a disk drive that is accessible to the local computer through a service such as Network Neighborhood, NFS, AFS, or DFS.

LPAR
logical partition.

machine
any type of data processing hardware that can be recognized by SAS IT Resource Management as a single unit. A machine can be a single physical machine such as a computer, a router, or a telephone switch. It can also be a logical entity such as a logical partition of an IBM zSeries system or a domain on the Solaris operating system. A computer in which multiple CPUs share the same memory could also be regarded as a single machine.

macro call
the programming code that invokes a macro.

macro source
another term for macro call.

macro variable
a variable that is part of the SAS macro programming language. The value of a macro variable is a string that remains constant until you change it. Macro variables are sometimes referred to as symbolic variables.

maintain a table
to apply changes from a table definition in a master data dictionary to a table definition in a performance data warehouse (PDB).

management information base
See MIB (management information base).

master data dictionary
in the SAS IT Resource Management program library, the master data definitions of tables and their associated variables, as well as override control statements, if any. Master copies of user-written definitions can be installed in the data dictionary, too.

MIB (management information base)
a virtual database for data that is gathered by one or more SNMP agents.

migrate a table
to add a table definition and the definitions of its associated variables from the master data dictionary to a performance data warehouse's (PDB's) data dictionary. When a table is migrated, any override %CPDDUTL control statements that are associated with the table definition or with the definitions of its associated variables in the master data dictionary are also applied to the PDB's data dictionary.

month-level data
performance data that is summarized for an entire month or for the current month-to-date. Month-level data is summarized from detail- level data.

MXG software
a collection of SAS programs (from Merrill Consultants) that operates on raw performance data that is created by z/OS, VM, and AS/400 operating systems and subsystems, and by other programs that log data in the same format.

network
a configuration of data processing devices and software that are connected for the purpose of information exchange.

operating environment
a computer, or a logical partition of a computer, and the resources (such as an operating system and other software and hardware) that are available to the computer or partition.

override control statements
one or more %CPDDUTL control statements that are automatically applied to a table definition in the active performance data warehouse (PDB) as the last stage of adding (migrating) a table definition and the definitions of its associated variables from the master data dictionary to the PDB's data dictionary. In the master data dictionary, the override statements, if any, are stored separately from the definitions of the table and its associated variables.

package
collector support that can be distributed to and incorporated into existing SAS IT Resource Management installations in order to provide additional, updated, or new functionality. See also collector support.

palette
a stored set of SAS and SAS/GRAPH options, such as colors and patterns, that can be associated with one or more report definitions.

parameter
a data item that is passed to a routine. In the SAS macro facility, a parameter is a constant or variable that is passed to the macro routine when the routine is invoked. See also keyword parameter, positional parameter.

PDB
See performance data warehouse.

PDB level
another term for level.

PDB options
characteristics of a specific performance data warehouse (other than data, table definitions, and variable definitions). PDB options include archive device, archive path, archive parameters, archive engine, active collector, Daylight Saving Time definition, Greenwich Mean Time deviation, start of week, work shift schedule, and definitions of work shift codes. On z/OS, you can also specify whether to build MXG views as a PDB option.

PDB's data dictionary
in a SAS IT Resource Management performance data warehouse (PDB), a SAS library whose name and libref are DICTLIB. The library contains metadata such as definitions of tables and variables and status information for the data that the tables contain.

performance data
information about how your IT system is doing its work. You use this information to determine whether your system is doing work efficiently and to determine how to change its load or its tunable parameters so that it will perform more efficiently.

performance data warehouse (PDB)
　　logically, a data warehouse that contains detailed and summarized performance data, as well as information (metadata) that is needed for managing the data. Physically, each PDB consists of a coordinated set of nine SAS libraries that contain performance data plus other information that is related to one or more IT services.

permanent SAS library
　　a SAS library that is not deleted when a SAS session ends, and which is therefore available to subsequent SAS sessions.

PGMLIB
　　the libref (library reference name) and library name for a SAS IT Resource Management program library that is on the local server host. See also RPGMLIB.

positional parameter
　　in a SAS macro, a parameter that is identified only by its position in the definition and invocation of the macro. Positional parameters must be separated by commas. If a positional parameter is not used, a comma must be inserted to indicate that the parameter is not being used. For example, if filename is the second positional parameter in a macro, you would either type the name of your data file in the second position (parm1, filename, parm3), or use a comma as a placeholder (parm1, , parm3).

process
　　to extract raw data from a data collector or data source, to optionally transform the data, and to load the data into a detail-level SAS library in a performance data warehouse (PDB).

process step
　　another term for process task.

process task
　　the task (or step) that processes the raw data into the detail level of a performance data warehouse (PDB).

program library
　　the SAS library in which most of the SAS IT Resource Management software resides. For example, the master data dictionary resides in the program library. See also PGMLIB, RPGMLIB.

protocol
　　a set of rules that govern data communications between computers, between computers and peripheral devices, and between software applications. TCP/IP, FTP, and HTTP are examples of protocols.

pull-down menu
　　the list of menu items or choices that appears when you choose an item from a menu bar or from another menu.

QuickStart Wizard
　　a wizard in the GUIs for the SAS IT Resource Management server. This wizard enables you to select a data collector and the type of log that you want to work with. The wizard creates a new performance data warehouse (PDB). The wizard also creates several batch jobs, including jobs that you can use to process data into the PDB, to reduce the data in the PDB, and to generate reports about the data.

reduce
　　another term for summarize.

reduce step
　　another term for reduce task.

reduce task
the task (or step) that reduces the data in a performance data warehouse (PDB) from the detail level to the summarized levels.

reduction level
another term for summary level.

regular variable
a variable whose values are acquired from the raw data and stored in the performance data warehouse (PDB). See also derived variable, formula variable.

remote PDB
a performance data warehouse (PDB) that SAS IT Resource Management can access by connecting to a remote host on which the PDB is available as a local PDB. See also local PDB.

remote server profile
a stored set of values that are used for connecting to a remote server host.

report
a graph or text that is generated by running a report definition.

report definition
a specification that is used for generating a report. A report definition includes information such as the table and level, the names of the variables, the report style, and other attributes.

report gallery
another term for gallery.

reporting tasks
tasks that generate or manage reports. Reporting tasks typically require 'read' access to a performance data warehouse (PDB), to a site library, or to a program library. See also administration tasks.

response time
the amount of time between the submission of a request for service and the user's receipt of the system's reply to that request.

restore
to recover the contents of a backup copy.

results
1) exceptions that are generated by running (evaluating) one or more rule definitions. The exceptions are written to the results data set. 2) reports (text, graphs, or both) that are generated by running one or more report definitions. The reports are written to a SAS catalog and then to a directory or PDS. See also exception.

retrieve
to recover some or all of the contents of an archive.

RPGMLIB
the libref (library reference name) for the program library on a remote server host. See also PGMLIB.

rule definition
a specification for detecting exceptions. A rule definition includes information such as the table and level, the expression, the message, and other attributes. See also exception.

SAS catalog
a SAS file that stores many different kinds of information in smaller units called catalog entries. A single SAS catalog can contain several different types of catalog

entries, such as graphs (.GRSEGs) and source code (.SOURCEs). See also SAS catalog entry.

SAS catalog entry
a separate storage unit within a SAS catalog. Each entry has an entry type that identifies its purpose to SAS. Some catalog entries contain system information such as key definitions. Other catalog entries contain application information such as window definitions, Help windows, formats, informats, macros, or graphics output.

SAS data set
a file whose contents are in one of the native SAS file formats. There are two types of SAS data sets: SAS data files and SAS data views. SAS data files contain data values in addition to descriptor information that is associated with the data. SAS data views contain only the descriptor information plus other information that is required for retrieving data values from other SAS data sets or from files whose contents are in other software vendors' file formats.

SAS data view
a type of SAS data set that retrieves data values from other files. A SAS data view contains only descriptor information such as the data types and lengths of the variables (columns), plus other information that is required for retrieving data values from other SAS data sets or from files that are stored in other software vendors' file formats. Both PROC SQL views and SAS/ACCESS views are considered to be SAS data views. SAS data views are of member type VIEW.

SAS datetime value
an integer that represents a date and a time in SAS software. The integer represents the number of seconds between midnight, January 1, 1960, and another specified date and time. For example, the SAS datetime value for 9:30 a.m., June 5, 2000, is 1275816600.

SAS IT Resource Management client
a computer on which SAS IT Resource Management software is installed with a client license. The client accesses PDBs on servers through a telecommunications protocol such as TCP/IP. Most client access is limited to 'readonly' and includes support for reporting, data visualization, and querying PDB definitions and status.

SAS IT Resource Management server
a computer on which SAS IT Resource Management software is installed with a server license. The server has 'write' access to the performance data warehouses (PDBs) and is used for processing and reducing data into one or more PDBs that are accessible through the host operating environment's file system. In addition, the server is used for administering PDBs and can also be used to schedule batch work, to create report and rule definitions, to generate reports, and to manage report and rule definitions and reports.

SAS library
a collection of one or more SAS files that are recognized by SAS and that are referenced and stored as a unit. Each file is a member of the library.

SAS procedure
a program that produces reports, manages files, or analyzes data and which is accessed with a PROC statement. Many procedures are included in SAS software.

scheduled job
in SAS IT Resource Management, one of a set of batch jobs that typically run at night and that process data into one or more performance data warehouses (PDBs), reduce data in the PDB(s), generate reports, and back up the PDB(s).

schema
the physical layout of logged or collected data.

server
See SAS IT Resource Management server.

setup
the task of customizing the collector and SAS IT Resource Management software in order to enable data from the collector to be processed into a performance data warehouse (PDB) and reduced in the PDB.

simple index variable
a variable for which an index is built. The index contains information about the values of the variable and about the observations that contain those values.

site
a group of performance data warehouses (PDBs) that are located on the same server host and which share the same default values, the same list of holidays, and the same work shift code for holidays. The shared information, known as site options or site properties, is stored in a site library. See also site library, site options.

site library
a SAS library in which site-specific information is stored. See also SITELIB, site.

site options
options whose values are intended as default values for all performance data warehouses (PDBs) for a specific site. These site options include archive device, archive path, archive parameters, archive engine, active collector, Daylight Saving Time definition, Greenwich Mean Time deviation, start of week, work shift schedule, and definitions of work shift codes. Site options can also include some site-level specifications that do not act as PDB defaults. These site options include the list of holidays and the work shift code for holidays. On z/OS, you can also specify whether to build MXG views as a site option.

SITELIB
the libref (library reference name) for the active site library. See also site library, active site library.

SNMP (Simple Network Management Protocol)
a protocol or standard that is used for exchanging network management information, which is information that enables the network to be managed.

SNMP agent
a software module that performs network management functions. See also SNMP (Simple Network Management Protocol).

staged data
raw data that has been read, optionally transformed, and either written to a SAS data set or made available through a SAS data view.

staging code
SAS statements that read raw data, transform it, and either write it to a SAS data set or make it available through a SAS data view.

suffix
in SAS IT Resource Management, the eighth character in the name of a variable at the day, week, month, or year level. The suffix indicates the type of statistic (for example, N means count, S means sum, and no suffix means average). Earlier characters in the name contain the name of the detail-level variable on which the statistic is based.

summarize
to calculate summary statistics in order to represent the values of variables in the detail level of a performance data warehouse (PDB). The summary statistics are stored in the day, week, month, and/or year levels of the PDB.

summary level
a performance data warehouse (PDB) level in which the performance data is summarized for a specific unit of time. The four summary levels are day, week, month, and year. The data at a particular summary level is stored in a corresponding library. For example, data that is summarized by day is stored in the DAY library of the PDB.

supplied collector support
collector support that is supplied with SAS IT Resource Management.

supplied format
a format that is shipped with SAS IT Resource Management.

supplied report definition
a report definition that is shipped with SAS IT Resource Management.

supplied rule definition
a rule definition that is shipped with SAS IT Resource Management.

supplied software
software that is shipped with SAS IT Resource Management.

supplied table definition
a table definition that is shipped with SAS IT Resource Management.

supplied variable definition
a variable definition that is shipped with SAS IT Resource Management.

table
the table definition, the status information about the data that is associated with the table definition, and the data itself. A table's data is contained in SAS data sets from each of the performance data warehouse's five data libraries: DETAIL, DAY, WEEK, MONTH, and YEAR. A table's metadata is contained in the performance data warehouse's DICTLIB library.

table definition
1) the attributes of a particular table. 2) the attributes of a table and the attributes of its associated variables. 3) the attributes of a table and of its associated variables, plus the status information about the data that is associated with the table and its variables. The meaning that is intended is explained where the term is used. See also table.

text report
a report that tabulates or prints data values or which produces low- resolution graphics. See also graph report.

throughput
the rate at which requests for work are serviced by a computer system.

time format
in SAS software, the instructions that tell SAS how to write numeric values as time values.

tool name
one of a pair of identifiers that specifies a particular set of collector support. The other identifier is collector name. For example, the pair of identifiers specifies which collector-specific staging code to run during the process task.

tuning
the process of tracking the performance of an IT service, comparing the performance to service objectives in order to identify bottlenecks and hardware problems, and then adjusting the service to eliminate the problems. Tuning also involves distributing work equitably to the available computer hardware in order to optimize the use of resources.

utilization
1) for resources that can be partially occupied (for example, memory), the fractional usage of a resource. 2) for resources that cannot be partially busy (for example, a CPU), the ratio of time that a resource is being used to the total elapsed time.

variable
a column in a SAS data set or in a SAS data view. The data values for each variable describe a single characteristic for all observations. See also derived variable, formula variable, regular variable.

view
logically, a set of observations that contains values for regular, derived, and formula variables. Physically, a view provides information about regular and derived variables and their locations, as well as source statements for calculating the values of formula variables. There is one view for each level of each table in a performance data warehouse (PDB).

Web browser
a software application such as Microsoft Internet Explorer or Netscape Navigator that provides access to information that is on the Internet or on your company's intranet.

Web gallery
another term for gallery.

Web-enabled report
a report that has attributes that enable it to be displayed by a Web browser.

week-level data
performance data that is summarized for an entire week or for the current week-to-date. Week-level data is summarized from detail-level data.

weight
a numerical coefficient that is assigned to an item and which indicates the relative importance of the item in a frequency distribution or population.

weighting variable
a variable whose values represent the weights for each observation.

workload
the amount of work a system is experiencing. Workload can be measured in terms such as characters per second, processes per second, or transactions per second.

year-level data
performance data that is summarized for an entire year or for the current year-to-date. Year-level data is summarized from detail-level data.

Index

Your Turn

If you have comments or suggestions about *Getting Started with SAS IT Resource Management 2.7*, please send them to us on a photocopy of this page, or send us electronic mail.

For comments about this book, please return the photocopy to

SAS Publishing
SAS Campus Drive
Cary, NC 27513
E-mail: **yourturn@sas.com**

For suggestions about the software, please return the photocopy to

SAS Institute Inc.
Technical Support Division
SAS Campus Drive
Cary, NC 27513
E-mail: **suggest@sas.com**