

SAS® IT Resource Management 3.21 Administrator's Guide

Second Edition



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2012. *SAS® IT Resource Management 3.21: Administrator's Guide, Second Edition*. Cary, NC: SAS Institute Inc.

SAS® IT Resource Management 3.21: Administrator's Guide. Second Edition

Copyright © 2012, SAS Institute Inc., Cary, NC, USA

All rights reserved. Produced in the United States of America.

For a hardcopy book: No part of this publication may be reproduced, stored in a retrieval system, or transmitted, in any form or by any means, electronic, mechanical, photocopying, or otherwise, without the prior written permission of the publisher, SAS Institute Inc.

For a Web download or e-book: Your use of this publication shall be governed by the terms established by the vendor at the time you acquire this publication.

U.S. Government Restricted Rights Notice: Use, duplication, or disclosure of this software and related documentation by the U.S. government is subject to the Agreement with SAS Institute and the restrictions set forth in FAR 52.227–19 Commercial Computer Software-Restricted Rights (June 1987).

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513.

1st electronic book, February 2012

SAS® Publishing provides a complete selection of books and electronic products to help customers use SAS software to its fullest potential. For more information about our e-books, e-learning products, CDs, and hard-copy books, visit the SAS Publishing Web site at

support.sas.com/publishing or call 1-800-727-3228.

SAS® and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.

Other brand and product names are registered trademarks or trademarks of their respective companies.

Contents

<i>About This Book</i>	<i>vii</i>
----------------------------------	------------

PART 1 Introduction 1

Chapter 1 • General Overview	3
Introduction to SAS IT Resource Management	3
Architecture and Basic Principles for Data Administrators	6
How to Locate Documentation	10
National Language Support (NLS)	12
Terminology Changes	13
Accessibility	14
Chapter 2 • Preparing to Work with the SAS IT Resource Management Client	15
Working with SAS Management Console	15
Working with SAS Enterprise Guide	20
Working with Metadata	27
Additional Requirements for z/OS	31
Chapter 3 • Working with the SAS IT Resource Management Client	33
Overview of the SAS IT Resource Management Client	33
Features of the SAS IT Resource Management Client	36
Working with Wizards	43
SAS IT Resource Management Transformations	45
Icons for SAS IT Resource Management Objects	46

PART 2 Accessing and Processing IT Data 49

Chapter 4 • IT Data Mart	51
About the IT Data Mart	51
Working with IT Data Marts	56
Importing, Exporting, and Promoting Metadata	60
Chapter 5 • Adapters	67
About Adapters	67
Working with Template Tables	75
Maintaining Template Tables and Staged Tables Based on Adapter Changes	84
Chapter 6 • Formulas	97
About Formulas	97
Working with Formulas	100
Chapter 7 • Staging the Data	105
About Staging the Data	106
Working with Staging Transformations	107
Working with Staged Tables	112

Columns in Staged Tables	124
Computed Columns in Staged Tables	126
Staging VMware Data Using Local Time	126
Understanding External Names of Staged Tables and Columns	127
Insufficient Disk Space	129
Chapter 8 • Aggregating the Data	131
About Aggregation Tables	132
Working with Aggregation Tables	140
Creating Aggregation Tables with the Summarized Aggregation Table Wizard	141
Creating an Aggregation Table with the Simple Aggregation Table Wizard	157
Creating Aggregation Tables with the Aggregations through Templates Wizard	159
Modifying Aggregation Tables	160
Working with Aggregation Transformation Filters	165
Working with Aggregation Transformations	170
Indexing an Aggregation Table	173
Chapter 9 • Adapter Setup Wizard	177
About the Adapter Setup Wizard	177
Using the Adapter Setup Wizard	179
What Does the Adapter Setup Wizard Create?	191
Domain Categories	204
Deploy and Run Jobs That the Adapter Setup Wizard Creates	205
Chapter 10 • Information Maps	207
About Information Maps	208
About Information Map Transformations	210
User-Specified Information Map Transformations	212
Information Map Objects That Are Generated by the Adapter Setup Wizard	219
Working with Information Map Filters	221
Using SAS Information Map Studio	226
Troubleshooting Information Maps Problems	229
Chapter 11 • User-Written Staging Code	233
User-Written Staging Transformations	233
What Are the Properties of the User-Written Staging Transformation?	234
Guidelines for Processing Data with User-Written Staging Transformations	235
Working with User-Written Staging Transformations	252
Using the Additional Generated Code	259
Chapter 12 • Jobs	263
Jobs That Process IT Data	263
Working with Jobs	268
Running Jobs	273
 PART 3 Appendixes 285	
Appendix 1 • SAS IT Resource Management Components	287
Appendix 2 • Data Sources Supported by SAS IT Resource Management Adapters	289
Data Sources Supported by SAS IT Resource Management Adapters	290
Supported Adapters	291
Accessing Raw Data for MXG	298
Accessing Raw Data for HP Perf Agent	298

Accessing Raw Data for DT Perf Sentry	300
Accessing Raw Data for BMC Perf Mgr	303
Accessing Raw Data for SAR	306
Accessing Raw Data for SNMP	308
Accessing Raw Data for Web Log	315
Accessing Raw Data for Database Adapters	315
Working with Custom Tables	328
Working with MXG Code to Stage Data	330
VMware vCenter Data Acquisition Jobs	351
Appendix 3 • Staging Parameters	355
Staging Parameters	356
Staging Parameters for Supported Adapters	362
Appendix 4 • Jobs That the Adapter Setup Wizard Creates for Domain Categories	381
Appendix 5 • Duplicate-Data Checking	391
Duplicate-Data Checking Overview	391
What Is Duplicate-Data Checking?	392
Control Data Sets for Duplicate-Data Checking	392
Content of Permanent Control Data Sets	393
Duplicate-Data-Checking Macros	394
Implementing Duplicate-Data-Checking Macros	394
Instructions for Implementing Duplicate-Data Checking	396
Examples	414
Appendix 6 • Adapter Key Metrics	417
What Are Key Metrics?	417
Reporting Areas or Domain Categories	418
Appendix 7 • Data Model	459
Overview of the SAS IT Resource Management Data Model	459
Data Model Adapter and Domain Category Specifications	470
Appendix 8 • Naming Standards	503
About Naming Standards for SAS IT Resource Management Objects	503
Naming Standards and Location of Objects	503
Appendix 9 • Macros	509
Introduction to the Macros in SAS IT Resource Management	510
%RMDELETE	511
%RMDUPCHK	513
%RMDUPDSN	516
%RMDUPINT	517
%RMDUPUPD	518
%RMMSBMCP	519
%RMMSDTPS	521
%RMSSSNMP	523
%RMPOINT	524
%RMRUNETL	529
Appendix 10 • Best Practices and Troubleshooting Tips	535
How to Back Up SAS IT Resource Management	536
Debugging Tips	539
Tips for Working with SAS IT Resource Management Client	541
Tips for Working with the IT Data Mart	541

Tips for Working with Aggregations	542
Tips for Working with Information Maps	543
Tips for Installing and Configuring SAS IT Resource Management	544
Handling Holidays	544
Performance Issues	549
Appendix 11 • Statistics	553
Index	555

About This Book

Audience

The purpose of the *SAS IT Resource Management: Administrator's Guide* is to present all topics that are related to system administration and data preparation in SAS IT Resource Management. For information about reporting topics, see *SAS IT Resource Management: Reporting Guide*.

SAS IT Resource Management: Administrator's Guide is designed for the following users:

- Data administrators are responsible for administering the IT data mart and setting up the jobs that stage and aggregate IT performance data so that report-ready data is available for generating reports. Data administrators typically deploy and schedule the batch production jobs that prepare and generate the IT performance reports. They often function as IT performance managers and capacity planners.

Data administrators typically use the documentation that is found in the *SAS IT Resource Management: Administrator's Guide*, although they usually refer to all SAS IT Resource Management documentation.

- Performance analysts are responsible for analyzing IT performance data that is managed by the data administrator, and designing and creating reports that communicate IT intelligence. These users analyze this data in order to best benefit the business and to improve the utilization, availability, and performance of IT resources and the IT enterprise. Performance analysts often function as capacity planners, system administrators, and business analysts.

Performance analysts typically use the documentation that is found in *SAS IT Resource Management: Reporting Guide*, although they might refer to all SAS IT Resource Management documentation.

- Information consumers are responsible for analyzing report data and making decisions based on that data. These users are business people who are interested in the overall performance aspects of IT at a company. They often function as system architects, IT managers, and IT executives. Information consumers use the reports that are generated by SAS IT Resource Management to support executive-level decision-making. They need to access the reports from the office as well as from remote locations.

Information consumers typically use the documentation that is found in *Introduction to SAS IT Resource Management* and in *SAS IT Resource Management: Reporting Guide*, especially the sections about SAS Web Report Studio, Gallery Manager, and interpreting reports.

Prerequisites

Here are the prerequisites for using SAS IT Resource Management:

- A user ID and password that is appropriate for the type of access that is needed to accomplish designated IT Resource Management functions.

Note: Data administrators use SAS Management Console to define the logins for users and groups of users, based on their roles and the SAS products that they need to work with.

- An operating environment that includes a supported client, middle tier, and server.
- **Depending on your role, access to the following software that is used by SAS IT Resource Management:**
 - SAS Data Integration Studio with plug-ins for SAS IT Resource Management
SAS Information Studio
 - SAS Enterprise Guide with add-ins for SAS IT Resource Management
 - SAS Add-In to Microsoft Office
 - Gallery Manager
 - SAS Web Report Studio
 - SAS BI Dashboard
 - SAS Information Delivery Portal
 - SAS Management Console

To log on to a SAS product, select **Programs** ⇒ **SAS**. From the list of software that displays, select the application that you want to work with. If prompted, enter your user ID and password.

Online Help is available for all SAS software. Click **Help** within the product in order to access it. In addition, documentation for SAS software is available at this Web location: <http://support.sas.com/documentation/index.html>. Use the A-Z index to locate the documentation that you want to access.

Part 1

Introduction

<i>Chapter 1</i>	
General Overview	3
<i>Chapter 2</i>	
Preparing to Work with the SAS IT Resource Management Client .	15
<i>Chapter 3</i>	
Working with the SAS IT Resource Management Client	33

Chapter 1

General Overview

Introduction to SAS IT Resource Management	3
What Is SAS IT Resource Management?	3
SAS IT Resource Management Functionality	4
Architecture and Basic Principles for Data Administrators	6
The Architecture of SAS IT Resource Management	6
The Basic Principles and Components of SAS IT Resource Management	7
How to Locate Documentation	10
Accessing SAS IT Resource Management Documentation	10
Accessing SAS IT Resource Management Online Help	12
National Language Support (NLS)	12
Terminology Changes	13
Accessibility	14

Introduction to SAS IT Resource Management

What Is SAS IT Resource Management?

Information Technology managers need to access, manage, integrate, manipulate, analyze, and share information about large quantities of performance data for many IT resources (such as hardware, operating system software, networks, Web servers, databases, and applications). As such, IT performance evaluation software requires technology that can work with performance data from these many different sources, summarize and analyze that data, and provide reports that permit quick and accurate analyses of a company's IT resources from current, historical, and forecasted perspectives. The SAS solution to this critical business challenge is SAS IT Resource Management (ITRM).

Note: The SAS IT Resource Management solution has more than one version of software. Throughout this document, the software for SAS IT Resource Management 2.7 or earlier contains the version number (for example, SAS IT Resource Management 2.7 or earlier). The software for SAS IT Resource Management 3.1.1 or later does not contain the version number and is referred to as SAS IT Resource Management.

SAS IT Resource Management provides many IT domain intelligence features. For example, it provides the following features:

- a comprehensive IT performance management and capacity planning data model for supported adapters
- the Adapter Setup wizard, which facilitates the efficient specification of the extract, transform, and load (ETL) jobs, analysis, and reporting processes for many IT resources
- tools for the creation of user-written adapters
- supplied report definitions that facilitate IT resource analysis

SAS IT Resource Management is built on SAS®. It provides the IT data mart structure, staging and aggregation code for specialized ETL processing of IT data sources, unique plug-ins to SAS Data Integration Studio and SAS Enterprise Guide, and a Web-based Gallery Manager application for the distribution and organization of IT intelligence reports. SAS IT Resource Management also supplies the IT Resource Management Gallery Manager application. Gallery Manager is a Web-based application that enables the many reports that are produced by the SAS IT Resource Management solution to be grouped, filtered, and managed in logical containers called galleries.

SAS IT Resource Management Functionality

SAS IT Resource Management can be used to resolve a variety of systems management challenges within an organization. It provides the functionality to do the following tasks:

- support (access, read, stage, aggregate, and report on) many popular network and systems management tools across hardware platforms, operating system environments, networks, Web servers, databases, and applications.
- establish IT data marts that contain the IT resource management data from data sources that the enterprise.
- create common references to common measurements that are available in disparate data sources. This functionality ensures that the IT data mart provides an independent set of measurements that can be analyzed across the IT enterprise.
- calculate standardized measurements, units, and statistics to populate the IT data mart with analysis and report-ready data.
- define Aggregation transformations that are specific to the time-based summarizations and statistical measurements that are necessary to accommodate a site's analysis and reporting needs.
- provide the ability to create graphical and textual reports that contain all the information that is needed to manage current day IT operations.
- use SAS IT domain knowledge for performance management, capacity planning, IT resource forecasting, prediction, peak period and seasonality analysis, time-based analysis, workload analysis, and enterprise IT performance summaries that complement existing IT utilization, availability, and performance analysis capabilities.
- provide a variety of report output mediums to accommodate the needs of the data administrators, performance analysts, and high-level information consumers.

It also takes advantage of SAS®9 and the enterprise-class solutions that SAS®9 provides, such as:

- SAS Data Integration Studio

The SAS IT Resource Management solution provides additional IT-specific features and functionality to the data management capabilities of SAS Data Integration Studio. These additional ITRM features and functions are seamlessly integrated with

those of SAS Data Integration Studio. The enhanced capabilities of SAS Data Integration Studio that are delivered by SAS IT Resource Management are commonly referred to as the SAS IT Resource Management client, throughout this document. They are labeled accordingly in the SAS IT Resource Management software.

The SAS Data Integration Studio solution provides flexible data integration services for most data sources. These services perform most of the data preparation and aggregation work that is needed to analyze and report on resource performance. You can use SAS Data Integration Studio to perform the following tasks:

- create and maintain the IT data marts that are the containers for the information maps, ETL jobs, tables, templates, data, and libraries that are created by or used by SAS IT Resource Management.
- create and run the processes that stage data in preparation for use by the Aggregation transformations.
- calculate new fields of data from the input data.
- create and run the processes that aggregate the data.

- SAS Statistics and Econometric Time Series

This solution provides a complete set of SAS statistical methods for IT data analysis.

- SAS Enterprise Guide

This solution provides comprehensive report definition and reporting services. Report definitions that are created in this solution can be run interactively within SAS Enterprise Guide, or they can be scheduled to run in batch mode.

- SAS Intelligence Platform

This platform uses the SAS Foundation technology, which includes the following software: Base SAS, SAS/GRAPH, SAS/STAT, the Output Delivery System (ODS), and other technologies. It also provides access to the following components:

- SAS Information Map Studio: This product provides the functionality to create business-oriented data sources that are appropriate for all users, expert to novice, when designing and creating reports.
- SAS Web Report Studio: This Web-based, interactive query and analysis tool simplifies and standardizes access to and reporting of performance evaluation information.
- SAS Metadata Repository: This repository of centralized metadata stores information about the objects that are created and used by SAS IT Resource Management, such as IT data marts, ETL jobs, information maps, and more.
- SAS Information Delivery Portal: This open Java portal is a customizable portal that provides corporate decision makers with easy access to the data that is most pertinent to them. It enables users to access data by means of reports and dashboards, such as key performance indicators (KPIs), filter reports, and more. SAS Information Delivery Portal enables the selective and secure dissemination of information throughout an organization.
- SAS Add-In for Microsoft Office: This product enables integration with commonly used Microsoft products such as Microsoft Word and Microsoft Excel.

Architecture and Basic Principles for Data Administrators

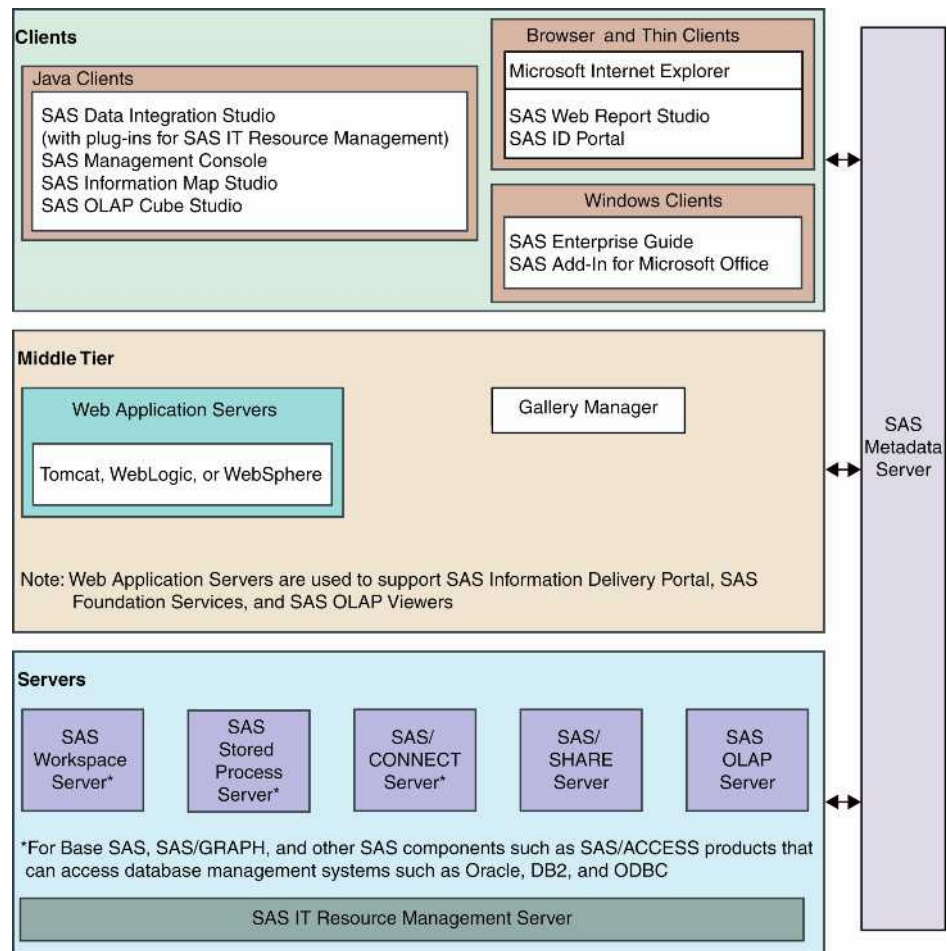
The Architecture of SAS IT Resource Management

SAS IT Resource Management uses data integration and business intelligence components that are provided by the SAS Intelligence Platform as well as additional software that was designed specifically for the processing of IT resource data. At installation time, users can deploy these components across multiple tiers. Using a multi-tier architecture to separate major software functions supports flexible processing schemes. The following functions can be supported on the tiers that SAS IT Resource Management uses:

- The client tier provides the interface between the user and SAS IT Resource Management. The client software enables you to perform data administration tasks, build reports, and view reports.
- The middle tier provides server support for Java applications. The Web-based Gallery Manager application runs on this tier.
- The server tier provides services that access data and perform background processing such as managing the metadata, executing stored processes, and performing resource aggregation calculations. In addition, the server tier provides the metadata for the files and tables that are input to and generated by the SAS IT Resource Management software.

With multi-tier processing, the tiers can be separated so that the processing can be done on multiple machines. The SAS IT Resource Management client must run on Windows. The SAS IT Resource Management middle tier and server can run in Windows, UNIX, or z/OS operating environments. Therefore, those software components for the server and the middle tiers can be installed on different machines according to the site requirements of an enterprise and the platforms that are supported by SAS for each tier of the SAS software architecture. For example, in a Windows environment, all of the processes that are needed by SAS IT Resource Management can be installed on a single PC. Other sites might prefer to install the client applications and middle tier on a PC and install the remaining components on the mainframe. Another option is to install the components needed for Web-based functions on a UNIX processor. The SAS representative can discuss these options and ensure the most appropriate configuration for a site.

The following diagram displays the architectural components that constitute the clients, middle tier, and servers that are used by SAS IT Resource Management.

Figure 1.1 Tiered Architecture of SAS IT Resource Management

Note: The multiple tiers shown in this diagram represent categories of software that perform similar functions. They do not necessarily represent separate machines or processors. In addition, your site might not need to use all of these software components. For example, if your site does not use the SAS Add-In for Microsoft Office, you do not need to install it.

The Basic Principles and Components of SAS IT Resource Management

Preliminary Tasks

SAS IT Resource Management has the flexibility to support customers whose usage of the product can range from the basic to the highly complex. Information consumers can work with specialized tools to access their IT data directly and generate ad hoc reports. Information consumers can use these reports to answer specific questions about the health and efficiency of the enterprise's IT resources. In all cases, the raw data for the resources that are being evaluated must be staged and aggregated. The resulting tables are input to the information maps and the reporting processes.

The following information describes the preliminary tasks that must be performed in order to use SAS IT Resource Management. After SAS IT Resource Management is installed and before the staging and aggregating processes can run, the data administrator uses SAS Management Console to perform the following setup tasks:

- Connect to the metadata server.
- Define new workspace servers.
- Add users and groups to the online SAS Metadata Repository.

Note: The installation process for SAS IT Resource Management provides the definitions that are needed for the supplied groups. No definitions are supplied for users, so the data administrator must add these definitions.

- Modify the default Access Control Template (ACT) for users or groups.
- Define SAS Content Server mapping for use by SAS IT Resource Management report packages. The SAS Content Server requires an HTTP server in order to be used by SAS Web Report Studio and other SAS Business Intelligence Platform applications. This server must be configured so that its content and its metadata are synchronized and are compatible with the file structure of SAS IT Resource Management.
- Define deployment locations for deploying jobs from SAS Data Integration Studio.

For more information about these tasks, see [“Working with SAS Management Console” on page 15](#).

Using IT Resource Management

After the preliminary tasks and initial setup are complete, you can use SAS IT Resource Management to process and evaluate IT resource data. SAS IT Resource Management provides the processes and the supporting technology that are required to regularly collect, aggregate, analyze, and report on the IT performance evaluation data that is vital to the health of an enterprise.

To start SAS IT Resource Management from the Windows **Start** menu, select **Programs** ⇒ **SAS** ⇒ **SAS IT Resource Management**.

In general, the processes that gather and monitor IT performance evaluation data consist of the following steps.

1. Collect the raw data about a resource.

IT performance data is information about IT resources, such as hardware, operating system software, networks, Web servers, databases, and applications. Raw (or unprocessed) data about the usage or performance of these resources is generated by the logging mechanisms that are inherent to IT resources or is created by the Enterprise Systems Management tools that are used to manage the IT infrastructure. The raw data is input to SAS IT Resource Management transformations that are supplied with SAS IT Resource Management software.

For information about how to generate and reference the raw data for supported adapters, see the various "Accessing Raw Data" topics in [“Data Sources Supported by SAS IT Resource Management Adapters” on page 290](#).

2. Process (stage) the raw data.

The raw data is processed (or staged) by adapters that are tailored to the data sources that they process. The adapter's staging code performs functions such as normalizing measurement units, generating computed columns from the raw data, and checking for duplicate data. Staging is performed by transformations that are set up in jobs. The job that actually executes the transformation can run interactively but is normally scheduled to run in batch mode, depending on the requirements of your enterprise.

A unique staging transformation is supplied for each adapter that is supported by SAS IT Resource Management. It contains the code and can access the associated

templates that are needed to process and load the raw data into staged tables. The staged tables can then be used as input to an Aggregation transformation or other SAS Data Integration Studio and SAS IT Resource Management transformations.

Information about these staging components is available in the following topics.

- For a list of the adapters that are supported by SAS IT Resource Management, see [“Supported Adapters” on page 68](#).
- For more information about user-written staging code, see [“User-Written Staging Transformations” on page 233](#).
- For more information about the staging transformations that are available for the supported adapters, see [“About Staging the Data” on page 106](#).

3. Aggregate the data in the staged tables.

After the raw data is staged, it can be input to the aggregation step. An aggregation is the act or process of grouping data, using an operation that produces a statistic such as a sum, average, minimum, or maximum. The term aggregation can also refer to the grouped data that results from such an operation. Aggregation transformations can generate summarized aggregated tables or simple aggregated tables.

- Simple aggregations do not undergo any summarizations; they simply append new data as it was read by the staging transformation to an existing aggregation table. (These tables are typically called detail aggregation tables.)
- Summarized aggregations read data from a staged or user-defined table and then categorize and aggregate data according to the specifications of the Aggregation transformation.

Jobs that contain Aggregation transformations are created and updated in SAS Data Integration Studio.

An Aggregation transformation can contain multiple aggregations that each generate an aggregation table of data. SAS IT Resource Management provides the ability to customize Aggregation transformations based on a site's requirements. An Aggregation transformation can be customized to perform aging and filtering of the aggregated data. It can perform calculations that create additional columns of data and join columns from different tables. It can support summarizations of data into any choice of time periods. The Aggregation transformation can also compute statistics on the data, calculate percent of change of a statistic over a specified period of time, and calculate the rank of a statistic or a class or ID column.

An unlimited number of aggregations can be created for each staged table. For example, the same data source can be input to both a daily aggregation and a weekly aggregation within the same or separate Aggregation transformations. For more information about working with Aggregation transformations and aggregation tables, see [“About Aggregation Tables” on page 132](#).

After the performance data has been aggregated, it is ready for the reporting processes.

4. Generate information maps.

For a supported adapter, SAS IT Resource Management generates transformations that create information maps that reference the tables of data that are generated by that adapter's Aggregation transformations. Information maps provide clearly labeled references for all data fields that are used to create and view reports. Information maps can be used in SAS Enterprise Guide and SAS Intelligence Platform applications such as SAS Web Report Studio in order to generate reports that provide domain intelligence about the adapters, such as CPU utilization, threshold analysis, and peak period analysis.

For more information about working with Information Map transformations to create and customize information maps, see “About Information Maps” on page 208.

5. Define and generate the reports.

SAS IT Resource Management generates tabular and graphical reports from report definitions that are supplied with the software, or created in SAS Enterprise Guide. Report definitions can be accessed by the Gallery transformation and, when the job that contains that Gallery transformation is run, the reports are generated and stored on the SAS Content Server. Report jobs are usually run in batch mode.

Report definitions can be run directly from SAS Enterprise Guide, and the resulting reports can be viewed interactively from the SAS Enterprise Guide Project window or from a Web browser. The reports can also be published to your Web server.

6. View the reports.

Reports that are generated by running the Gallery transformation are stored on the SAS Content Server where they can be accessed and managed by Gallery Manager. Gallery Manager enables information consumers to organize their reports into galleries and to set accessibility controls on those galleries. Galleries can be marked as private (for an individual's use) or public (for use by all Gallery Manager users.) Galleries can also be filtered so that they display only a subset of the available reports.

Reports can also be viewed on the Web through SAS Information Delivery Portal.

For more information about the reporting capabilities that are available using SAS IT Resource Management data, see *SAS IT Resource Management: Reporting Guide*.

Information consumers and performance analysts can also use these additional processes:

- Generate and view ad hoc reports.

SAS Web Report Studio is a Web-based interface that can use information maps to select and report on data. Information maps are generated by Information Map transformations.

- Use Microsoft Office products.

SAS Add-In for Microsoft Office enables SAS functionality to be accessed directly from the menus and toolbars of Microsoft Office word processing and spreadsheet products.

Performance analysts and information consumers who frequently work with Microsoft Office products such as Microsoft Word and Microsoft Excel can use SAS to access, analyze, and report on IT performance data that is available in information maps, and then distribute the results to other people in the enterprise.

How to Locate Documentation

Accessing SAS IT Resource Management Documentation

Documentation about SAS IT Resource Management is available in the form of Web-based documentation, and online Help.

To locate the SAS IT Resource Management documentation, use the Products Index at <http://support.sas.com/documentation/index.html>. The following

documentation can be accessed from the SAS IT Resource Management documentation Web page:

- *What's New in SAS IT Resource Management*

This document provides a brief description of the new features that are included in this release of the SAS IT Resource Management software and documentation.

- *Introduction to SAS IT Resource Management*

This document provides information about the extensive domain intelligence that is available with SAS IT Resource Management.

The audience for this introductory document is anyone who works with SAS IT Resource Management or who uses the reports that are generated by the software.

- *SAS IT Resource Management: Administrator's Guide*

This document contains detailed information about the tasks that are required to set up and manage the IT data mart. It also describes how to set up the IT Resource Management adapters so that they can process the IT resource data that is stored in the IT data mart. This guide provides information about how to set up and maintain IT data marts, work with simple and summarized aggregations, information maps, and the wizards, especially the Adapter Setup wizard, that facilitate the specification of the ETL processes required for administering a site's IT resource data.

This document provides information about the other functional components of the software, such as the SAS Metadata Repository (SMR) and the like.

The audience for this document is SAS IT Resource Management data administrators or capacity planners.

- *SAS IT Resource Management: Reporting Guide*

This document contains detailed information about the tasks that are required to generate, view, and manage reports. It describes how to work with the Gallery task in SAS Enterprise Guide, the Gallery transformation, and Gallery Manager application. It also provides information about the other SAS products that support the reporting processes, such as SAS Web Report Studio, SAS Information Delivery Portal, and SAS Add-in for Microsoft Office.

A "Report Conversion" appendix provides a description of the steps that are performed in order to re-create reports that were generated using SAS IT Resource Management 2.7 report macros. These steps use SAS Enterprise Guide, which is included in SAS IT Resource Management.

- *Gallery Manager: User's Guide*

This document contains detailed information about the tasks that are required to create, view, and manage galleries and to view the reports that are included in those galleries. (This information is also included in the *SAS IT Resource Management: Reporting Guide*.)

- *SAS IT Resource Management: Migration Documentation*

This document provides a description of the steps that are performed in order to move a site's processing from the earlier versions of the software to the current version.

- *SAS IT Resource Management Glossary*

This glossary provides definitions for term that are used in SAS IT Resource Management.

Documentation is available for the business intelligence and data integration components of the SAS Intelligence Platform and all the SAS products that are referenced in this document. To access this documentation, navigate to the Web site at this location:
<http://support.sas.com/documentation/index.html>.

Accessing SAS IT Resource Management Online Help

Online Help is available for anyone who performs tasks with the user interface.

Note: The Help that is provided with this release of the SAS IT Resource Management software briefly describes the fields on the windows that are specific to IT Resource Management. In addition, the online Help displays a collapsed, HTML version of the *SAS IT Resource Management: Administrator's Guide*, which includes concepts and descriptions of tasks that pertain to managing your IT resources. To view the hierarchical list of topics in this document, click the icon that contains a plus sign.

To display the field-based Help window for an active window or tab, click **Help** or press the F1 key.

To search for topics about concepts or features that are identified by specific words, such as "information maps," click the **Search** tab in the Help window. Enter the text to be found and press the ENTER key.

Online Help is available for all SAS technology products on which SAS IT Resource Management is built.

National Language Support (NLS)

National Language Support (NLS) is a set of features that enable a software product to function properly in every global market in which the product is sold. Typically, software that is written in the English language works well for users who use English and use data that is formatted using the conventions that are observed in the United States. However, without NLS, these products might not work well for users in other regions of the world. NLS in SAS enables users in regions such as Asia and Europe to process data successfully in their native languages and environments.

SAS IT Resource Management uses NLS formats for dates and numbers in template table columns, staged table columns, and aggregation table columns. If a column uses NLS formatting, the value in the **Format** column of the table begins with "NL" such as **NLNUM16.0**. These NLS values in the **Format** column of the tables use the following naming convention:

- **NLNUM** indicates that the data value is numeric.
- **NLDATE** indicates that the data value is a date (without the time component).
- **NLDATM** indicates that the data value is a date (with the time component).
- The number before the period indicates the maximum number of characters (including numerals, commas, and dots) that the data value can have.
- The number following the period indicates the number of decimal places that the data value can have.

For example, if a table column has the value **NLNUM16.2** in the **Format** column of the table, then you know that the data value is numeric, it has a maximum of 16 characters (including numerals, commas, and dots), and it has two decimal places. In this case, NLS enables the same data value to be presented in different ways (such as 1,000,000,123.75

in the United States or 1.000.000.123,75 in some European regions) based on the system locale.

Dates and times have many representations also, depending on the conventions that are accepted in a culture. The month might be represented as a number or as a name. The name might be fully spelled or abbreviated. The order of the month, day, and year might differ according to locale. Likewise, time can be represented in one English-speaking country or region by using the 12-hour notation, while other English speakers expect time values to be formatted using the 24-hour notation.

In order to have NLS present data in the format that represents the traditional form that your users might expect, you can specify the locale setting when invoking SAS. A locale reflects the language, local conventions such as data formatting, and culture for a geographical region. Local conventions might include specific formatting rules for dates, times, and numbers, and a currency symbol for the country or region. You can use the `LOCALE=` system option to specify the locale of the SAS session at SAS invocation.

Note: For more information about how to set the locale for your system, see the *SAS National Language Support (NLS): User's Guide*.

Terminology Changes

The versions of SAS IT Resource Management 3.1 and later use new terminology to describe similar tasks, objects, or features that you might have used when working with earlier versions of the software. The following table provides some of the old terms that were used before this version of SAS IT Resource Management and the new terms that describe the same or similar tasks or objects.

Table 1.1 Terminology Changes for SAS IT Resource Management 3.1, 3.1.1, 3.2, and 3.21

Old Term or Phrase	New Term or Phrase
PDB	IT data mart
collectors	adapters
QuickStart wizard	Adapter Setup wizard
detail level data	simple aggregation data
summary level data	summarized aggregation data
process macro	staging transformation
reduction macro	Aggregation transformation

Note: The terms and phrases in the table show the change in terminology for similar objects and concepts. These terms do not represent an exact one-to-one correlation. For example, the Adapter Setup wizard provides more functionality and features than the QuickStart wizard. However, the general concepts and objectives of the two wizards are similar.

Accessibility

SAS IT Resource Management includes the following accessibility and compatibility features that improve usability of the product for users with disabilities. These features are related to accessibility standards for electronic information technology that were adopted by the U.S. Government under Section 508 of the U.S. Rehabilitation Act of 1973, as amended.

SAS IT Resource Management addresses two areas of 508 compliance:

- Web accessibility compliance applies to browser-based content and Web-based applications.
- Software accessibility compliance applies to software. In SAS IT Resource Management, software accessibility compliance applies to the SAS Data Integration Studio user interface. SAS IT Resource Management uses the functionality of SAS Data Integration Studio and has the same accessibility features as that product. For information about these features, see the *SAS Data Integration Studio: User's Guide*. This documentation is located at this address: <http://support.sas.com/documentation/index.html>.

If you have questions or concerns about the accessibility of SAS products, send e-mail to accessibility@sas.com.

Chapter 2

Preparing to Work with the SAS IT Resource Management Client

Working with SAS Management Console	15
Introduction to SAS Management Console	15
Adding Users and Groups of Users	16
Adding Locations for the Deployment of Jobs	18
Working with SAS Enterprise Guide	20
About Working with SAS Enterprise Guide	20
Setting Up SAS Enterprise Guide for SAS IT Resource Management Data	20
Migrating SAS Enterprise Guide Projects to Work with SAS IT Resource Management 3.2 and Later	22
Working with Metadata	27
About Metadata	27
Supplied Metadata	28
System Authentication When Logging On to the SAS Metadata Repository	28
Backing Up and Restoring the SAS Metadata Repository	29
Migrating from an Earlier Version of SAS IT Resource Management	29
Tools for Exploring SAS IT Resource Management Metadata	30
Metadata Server Memory Issues	31
Additional Requirements for z/OS	31

Working with SAS Management Console

Introduction to SAS Management Console

After successfully executing the installation and configuration instructions that are delivered with SAS IT Resource Management, the data administrator must perform the setup tasks that are required by SAS IT Resource Management. These tasks ensure that performance analysts and information consumers have the permissions and tools that they need in order to work with their IT resource data. The data administrator performs these tasks with SAS Management Console and SAS Enterprise Guide.

Note: Unlike the previous versions of SAS IT Resource Management (versions 3.1 and 3.1.1), no specific customizations are needed for SAS IT Resource Management 3.2 and later.

SAS Management Console provides a single point of control for managing resources on all platforms that are supported by SAS. By invoking SAS Management Console, the data administrator has access to the many management functions that it provides. For information about these functions, see the *SAS Management Console: User's Guide* at

<http://support.sas.com/documentation/index.html>. Use the A-Z index to SAS Product Documentation to locate the documentation for SAS Management Console.

SAS Management Console enables SAS IT Resource Management data administrators to perform the following tasks:

- Add users and groups to the SAS Metadata Repository. By creating user groups, the data administrator can set permissions that provide security for data and reports that are registered in the SAS Metadata Repository. These groups govern how group members can access the various components of the system.

Note: The **User Manager** component of SAS Management Console controls this function.

- Add users and groups to the Access Control Template (ACT) table. When you add users or user groups to the ACT table, you can also set permissions that govern the tasks that those users or user groups can perform. The permissions govern the resources that have been registered in the SAS Metadata Repository.

Note: The **Authorization Manager** component of SAS Management Console controls this function.

- Add locations (or directories) for the deployment of jobs that are created in SAS IT Resource Management. After users define their jobs, they can submit those jobs for immediate execution or deploy them for scheduling in SAS Management Console. After a job is deployed, you can schedule the job to run in batch mode at a specified date and time.

Note: The **Schedule Manager** component of SAS Management Console controls the function of specifying a job deployment directory.

To access SAS Management Console, from your Windows desktop, select **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS Management Console**.

Adding Users and Groups of Users

About SAS IT Resource Management User Groups

The installation process of SAS IT Resource Management establishes two user groups:

- IT Resource Management Administrators

IT Resource Management Administrators can access all folders and galleries in Gallery Manager.

- IT Resource Management Users

IT Resource Management Users can access their own private folders and any public folders and galleries. In addition, they can access only their own private folders in Gallery Manager. (Public folders can be accessed by anyone.)

Note: Users who do not belong to either group have Read-Only access to public folders in Gallery Manager. In addition, Gallery Manager honors only the IT Resource Management Administrators and IT Resource Management Users groups. It does not honor any other groups that might be established.

SAS Management Console enables you to add new or existing users to a SAS IT Resource Management user group. To invoke SAS Management Console, perform the following steps:

1. Start SAS Management Console by selecting **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS Management Console**.

Connect to SAS Management Console by using the unrestricted user ID.

2. In the **Repository** field at the top of the **Plug-ins** tab, select the Foundation repository from the drop-down list.
3. With SAS Management Console, you can perform the following tasks:
 - “Add a New User to a SAS IT Resource Management User Group” on page 17.
 - “Add an Existing User to a SAS IT Resource Management User Group” on page 17.
 - “Set Permissions for Users” on page 18.

After adding a new or existing user to one of the IT Resource Management groups, set the login and then the permissions for that user. (A login is a copy of information about the external account for each application authentication domain.)

Add a New User to a SAS IT Resource Management User Group

To add a new user to one of the two groups that are provided by SAS IT Resource Management, perform the following steps:

1. Right-click the **User Manager** component of SAS Management Console. From the drop-down list, select **New** ⇒ **User**.

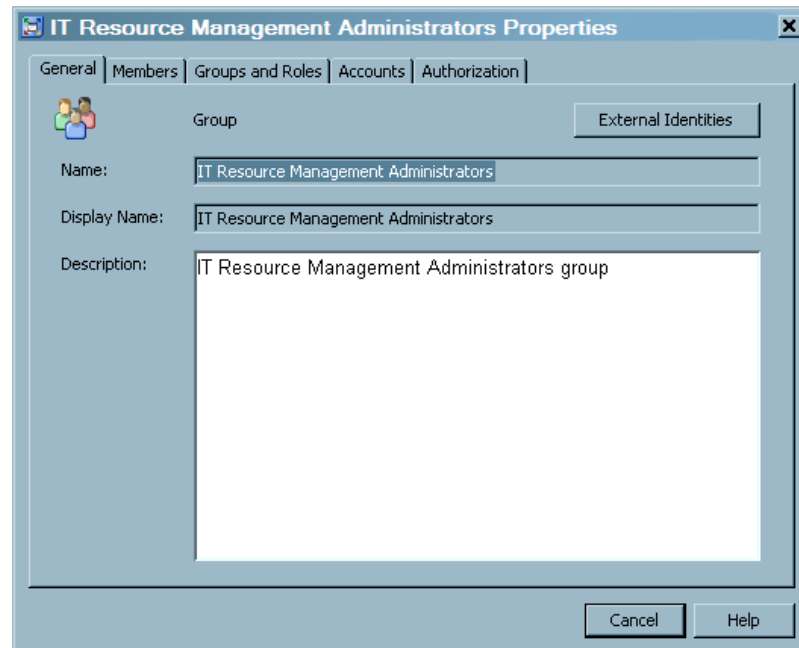
The **General** tab of the New User Properties dialog box appears. Enter the name and other identifying information for the user on this dialog box.
2. Click the **Groups and Roles** tab. From the **Available Groups and Roles** column, select the group (that is, IT Resource Management Administrators or IT Resource Management Users) into which you want to place the new user. Click the arrow to add the selected group to **Member of** column.
3. Define a login for this new user.

Click the **Accounts** tab. Click **New** to display the New Login Properties dialog box where you can add login information for the new user. Ensure that the **User ID** and **Password** fields are entered. (DefaultAuth is the supplied entry that contains the login for the SASITRM server.)
4. Click **OK** to save your changes.

Add an Existing User to a SAS IT Resource Management User Group

To add an existing user to one of the two groups that are provided by SAS IT Resource Management, perform the following steps:

1. From the list of registered users in the **User, Group, or Role** column, right-click the **IT Resource Management Administrators** or **IT Resource Management Users** and select **Properties**.

Display 2.1 IT Resource Management Administrators Properties Dialog Box

The dialog box contains five tabs: **General**, **Members**, **Groups and Roles**, **Accounts**, and **Authorization**. The Properties dialog box enables you to view and modify information about the selected user or user group.

2. Click the **Members** tab to display a list of the **Available Identities** and the **Current Members**. To add a user to the **Current Members** of selected group, highlight that user in the **Available Identities** column. Use the arrow to move the selected user to the **Current Members** column.
3. For each selected member, make sure that user has a login.

To set the login for a user, click **Properties**. On the **Accounts** tab, ensure that the **User ID** and **Password** fields are entered. (DefaultAuth is the supplied entry that contains the login for the SASITRM server.)

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

Set Permissions for Users

Permissions can be set in the ACT table in the **Authorization Manager** component of SAS Management Console. The ACT consists of a list of users and groups and indicates, for each user or group, whether permissions are granted or denied. An entry in the table can be applied to multiple resources. Make sure that the users have the appropriate permissions for the Foundation repository.

Note: For more information, see the *SAS Intelligence Platform: Security Administration Guide*.

Adding Locations for the Deployment of Jobs

In order to run a job in batch mode, a job must be deployed into a directory on the application server. The location of this directory can be defined in SAS Management Console. New deployment directories can also be defined in the SAS IT Resource Management Client, while you are deploying a job. To specify the location, perform the following steps:

1. If SAS Management Console is not running, start it by selecting **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS Management Console**.

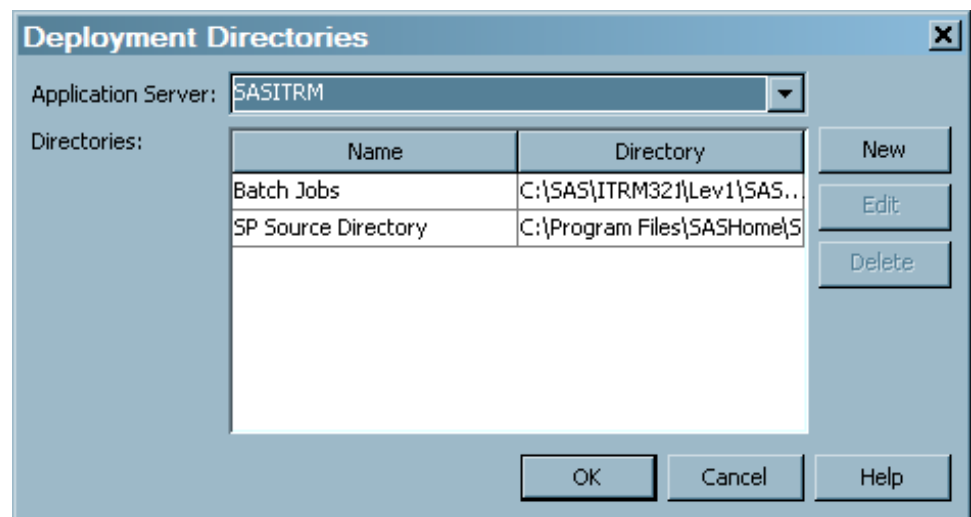
Note: You must have the authority to use the Scheduler Manager plug-in in SAS Management Console. Your system administrator can grant this authority, if necessary.

2. Click the **Plug-ins** tab, if necessary, to display plug-ins in the left panel of SAS Management Console. In the **Repository** field at the top of this panel, select the Foundation repository from the drop-down list.
3. Select the **Schedule Manager** component of SAS Management Console.

Note: If the **Schedule Manager** component is not displayed in the list of plug-ins, your user ID is not defined to have access to Schedule Manager. Users can be granted access to Schedule Manager by administrative users.

4. From the menu bar, select **Actions**. From the drop-down list, select **Deployment Directories**. The Deployment Directories dialog box appears.

Display 2.2 Deployment Directories Dialog Box



5. In the **Application Server** field, you can specify the application server that you want to use for deploying jobs. Use the drop-down list to show the application servers that are available, and click to select a server. (SASITRM is the application server that is configured by default when SAS IT Resource Management is installed.)
6. In the **Directories** section of the Deployment Directories dialog box, you can specify one or more directories where the generated code for deployed jobs is to be stored. The server that is selected in the **Application Server** field must be able to resolve the paths to these directories.

SAS IT Resource Management supplies a **Batch Jobs** directory. The path to this supplied directory on the SASITRM server is: **SASEnvironment\SASCode\Jobs**.

In the Deployment Directories dialog box, you can perform the following tasks:

- Define a new directory. To do so, click **New**.
- Update a directory. To do so, select the directory and click **Edit**.
- Remove the metadata for a directory. To do so, select the directory and click **Delete**.

If you are defining locations in z/OS traditional file system locations, you can specify the name of the directory as the full name of a partitioned data set.

For more detailed information about these tasks, click **Help**.

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

Working with SAS Enterprise Guide

About Working with SAS Enterprise Guide


SAS Enterprise Guide is one of the primary reporting vehicles for SAS IT Resource Management. The data administrator uses the functions of SAS Enterprise Guide in order to perform the following tasks:

- Set up SAS Enterprise Guide so that it can work with SAS IT Resource Management data.
- Migrate SAS Enterprise Guide projects so that they can work with your SAS IT Resource Management 3.2 and 3.21 IT data marts, in which you created the report data.


Setting Up SAS Enterprise Guide for SAS IT Resource Management Data

In order to work with the data from SAS IT Resource Management, SAS Enterprise Guide needs to access the SAS Metadata Repository. To access a SAS Metadata Repository for use by SAS Enterprise Guide, perform the following steps:

1. Invoke SAS Enterprise Guide. To do so, from the **Start** menu, select **Programs** ⇒ **SAS** ⇒ **Enterprise Guide**.
2. From the menu bar, select **Tools** ⇒ **Options**. The Options dialog box appears. From the list in the left pane, select **Administration**.

Display 2.3 Administration Window of the Options Dialog Box

3. If you want to create a profile or change your profile, click **Modify**, which opens the Connections window. Select the profile that you want to change and click **Modify**.
4. The Modify Profile window appears and enables you to create or revise the profile. Enter or revise the following information.
 - a. The name of the profile is required. The description is optional.
 - b. Specify whether the machine is remote or local by clicking the appropriate button.
 - c. Enter the machine and port number of the server that you want to connect to. (If you selected **Local**, the machine name is automatically set to localhost.)
 - d. Enter the user ID and password for the person who can access this server.
 - e. Click **Save** to save all specified values and close the window, or **Cancel** to discard all values and close the window.
5. To connect to the server, select the appropriate profile and click **Set active**. The active server is the source for all resource definitions (such as servers and libraries).
If the Credentials Required dialog box appears, enter your user ID and password and click **OK**.

The selected repository is identified by the active repository icon: .

6. Click **Close** to return to the Administration window.
7. Click **OK** to return to the main SAS Enterprise Guide workspace.

For information, see *Administering SAS Enterprise Guide*. To locate this documentation, navigate to <http://support.sas.com/documentation/onlinedoc/index.html>. From the list of products, select the **SAS Enterprise Guide**.

Migrating SAS Enterprise Guide Projects to Work with SAS IT Resource Management 3.2 and Later

About the Migration Wizard

The Migration Wizard for SAS Enterprise Guide and the SAS Add-in for Microsoft Office enables you to report projects that are created with earlier versions of SAS IT Resource Management and SAS Enterprise Guide. The upgraded version can then be used with your current version of SAS IT Resource Management and SAS Enterprise Guide.

SAS Enterprise Guide projects supplied by SAS IT Resource Management, and typically those created for use with SAS IT Resource Management, use information maps in the report projects. Therefore, you must change the map paths of the supplied information maps to point to the information maps that you created with the Adapter Setup wizard. The easiest way to change the map paths is to use the Migration Wizard for SAS Enterprise Guide and the SAS Add-in for Microsoft Office.

Prerequisites to Running the Migration Wizard

Before you run the Migration Wizard, you must run the staging, aggregation, and information map jobs that create the information maps for the adapter and domain category that you want to work with.

If you installed the supplied SAS Enterprise Guide projects, they can be accessed by navigating to the location where you installed SAS IT Resource Management. The projects are located at

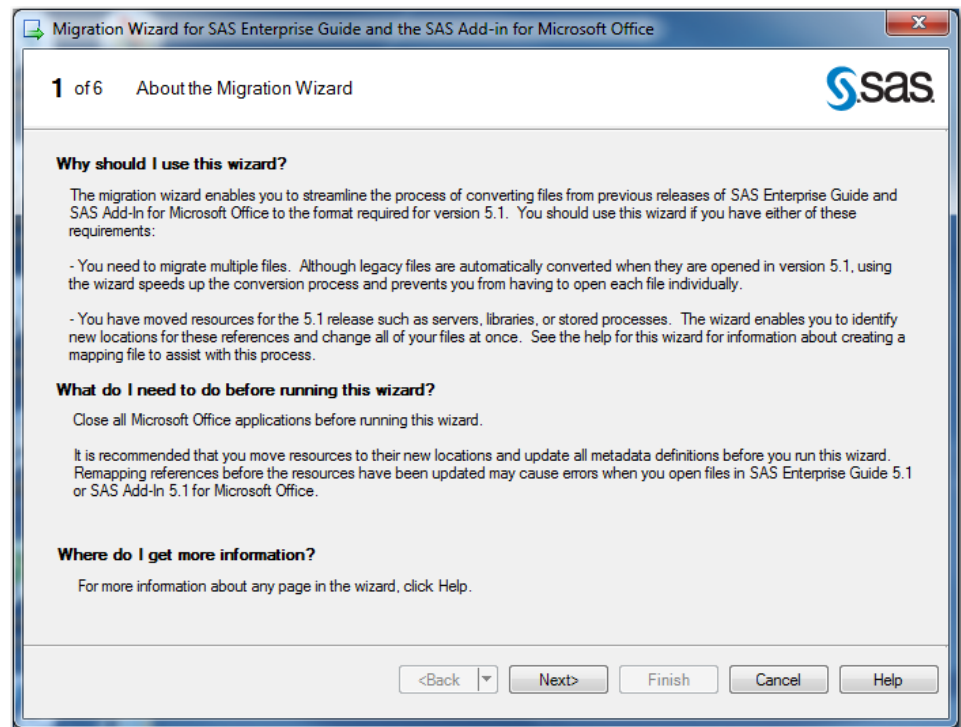
SASITResourceManagementSASEnterpriseGuideComponents\3.21. For example, if you installed SAS IT Resource Management on a 64-bit machine, the projects might be located at **C:\Program Files\SASHome\amd64\SASITResourceManagementSASEnterpriseGuideComponents\3.21**.

TIP Make sure the projects are not read-only projects. To do this, right-click each project and select **Properties**. On the **General** tab, make sure that the **Attributes** box is not checked. Then click **OK**.

TIP Make a copy of the SAS Enterprise Guide projects and migrate only the copied projects. The original projects should be kept as a backup.

Running the Migration Wizard

To run the Migration Wizard, navigate to the location where the current version of SAS Enterprise Guide is installed. Then select MigrationWizard.exe. Click **Enter** to invoke the Migration Wizard. The first page of the Migration Wizard appears. This page contains information about why you should use this wizard and what you should do before you run it.

Display 2.4 Page 1 of the SAS Enterprise Guide Migration Wizard

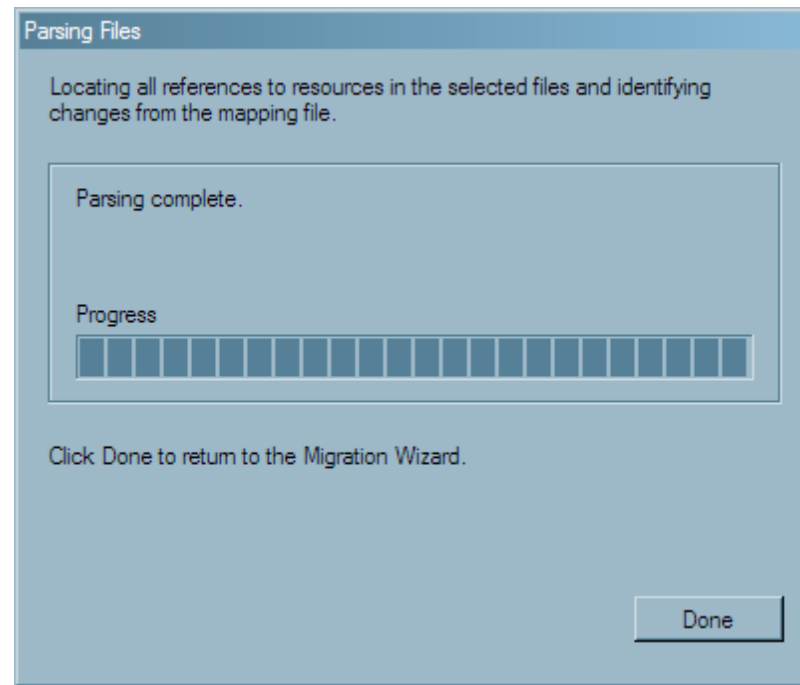
On page 2 of the wizard, verify the connection profile of the application server where your information maps are located. You can click **Modify** to change your connection profile.

On page 3 of the wizard, select the files that you want to migrate. To do so, perform the following steps:

1. To facilitate the selection of files, from the drop-down list in the **Show** field, select **SAS Enterprise Guide Projects (*.egg)**. This action causes only SAS Enterprise Guide projects to be displayed in the **Available files** column.
2. In the **Available files** column, navigate to the location of the SAS Enterprise Guide projects that you want to migrate.
3. Check the boxes next to the projects to be migrated.
4. Click the right arrow to move the projects to the **Selected files** column.

On page 4 of the wizard, information is displayed that explains the next steps of the migration process. Click **Next** to run the process that parses the files. The progress of the wizard is displayed as it locates all the references to the projects that you selected.


Display 2.5 Parsing Files Information Box




Click **Done** when the progress indicator shows that this task is complete.

On page 5 of the wizard, you must enter the new mapping information for the selected projects. To do so, perform the following steps:

1. Use the **Show** field to shorten the number of entries on this page. From the drop-down list, select **Global references only**. (This setting removes the information maps from this page, which significantly simplifies the contents of this page.)

Display 2.6 Selecting Only Global References

- Expand the page to the full width of your screen. This action enables you to have space to see the complete path information for the current and new mappings.
- If necessary, resize the **Current Mapping** column width so that you can see the complete path. This path is the path that was used in the SAS Enterprise Guide project when it was last saved. It identifies the adapter and domain category for which the project was developed. As shown in the following display, a typical path is quite long (for example, **/Shared Data/SAS IT Resource Management/IT Data Marts/VMwareForReporting2/VMware vCenter 1/Domain Categories/Guest System/**).

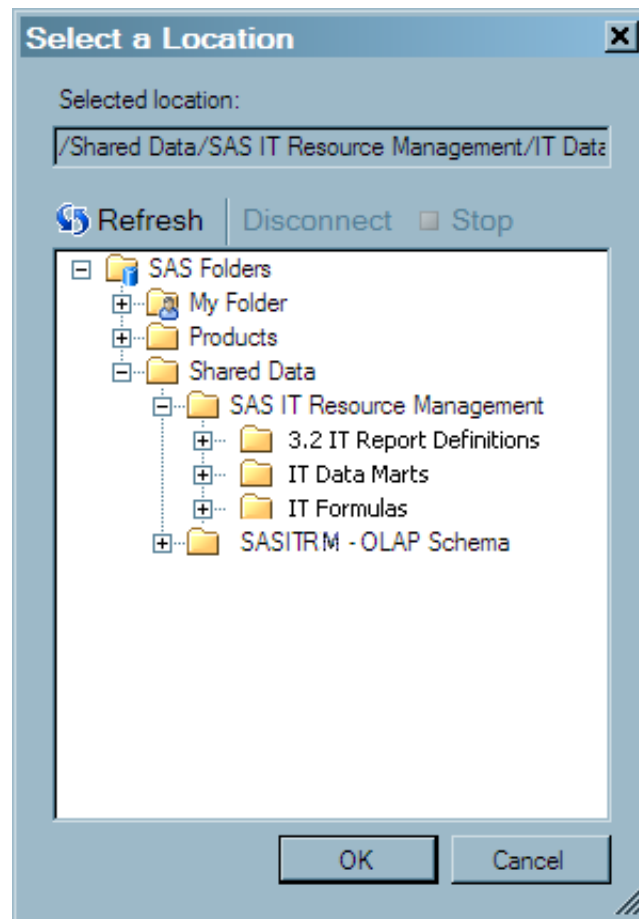
Display 2.7 Expanded Current Mapping Column

- If you want to globally change all references to the server, click the ellipsis button (...) in the **New Mapping** column. The Change Resource dialog box appears. Use the **Browse** button to display the list of servers from which you can select the **Server**

name that you want to use in your paths. If you want to apply this change to all the object references that contain this value, check the corresponding box.

5. You must enter the mapping information for your new IT data mart in the **New Mapping** column for each of the project folders that are shown on this page. To do so, perform the following steps:
 - a. Select the folder and click the ellipsis button (...) in the **New Mapping** column. The Change Resource dialog box appears.
 - b. Use the **Browse** button to display the list of folders to which the projects can be migrated. Expand the SAS Folders tree until you locate the IT data mart from which you ran the Adapter Setup wizard for these projects.

Display 2.8 Locating the IT Data Mart



- c. Continue expanding the folders within the IT data mart until you find the location of the adapter and the domain category that corresponds to the path in the **Current Mapping** field. In the Select a Location dialog box, the **Selected location** field displays the complete path for the new location.
 - d. Make sure that the box is checked that specifies to apply this change to all object references that contain this value.
6. When you have entered the new mapping information for all the SAS Enterprise Guide project folders, click **Next**.
7. On the Begin the migration process page, you can also enter the path to your Migration Record File. You can also change the path to the Log File, specify

whether to create automatic backups of the original files, and specify whether to verify the existence of all resources that were changed.

TIP For best results, specify that the original files should be backed up automatically.

8. Click **Finish** > to perform the migration. The Migrating files window displays the progress of your migration. When the migration is complete, the status of the processing is displayed. If a processing error occurred, you can view the log file. Otherwise, click **Done** to exit from the Migration Wizard. The jobs that contain the projects that were successfully migrated can be opened in the IT data mart and run.

Working with Metadata

About Metadata

Metadata is structured data that describes the location and structure of enterprise data. During the installation of SAS IT Resource Management, metadata is created and stored in the SAS Metadata Repository. The metadata for SAS IT Resource Management consists of the following information:

- where SAS servers are deployed and how they are configured
- the location of the physical SAS data and how to access it
- the attributes of the data, such as types, formats, and dimensions of the tables that contain the data
- the specifications for processes and jobs, such as ETL processes
- security issues, such as credentials, user identities, groups, and roles

The metadata for SAS IT Resource Management is created, maintained, and used by the following applications:

- **SAS Management Console:** creates and maintains metadata that describes the attributes of SAS servers that are used by SAS IT Resource Management. For example, it describes where the servers are deployed and how they are configured. The metadata also maintains lists of authorized users, their permissions, their account information, and more.
- **SAS IT Resource Management Client:** creates and maintains metadata that describes the attributes of the data that is used by SAS IT Resource Management. For example, it describes the tables of staged or aggregated data, the definitions of the columns within those tables, the jobs, the libraries, and the transformations that make up the data management component of SAS IT Resource Management.
- **SAS Enterprise Guide:** uses metadata that describes the servers, users, libraries, information maps, and information map filters for reporting with SAS Enterprise Guide.
- **SAS Web Report Studio:** creates and maintains metadata for Web-based reports, and uses other metadata created and maintained by other SAS clients.

Supplied Metadata

In addition to the metadata that is described in the previous topic, SAS IT Resource Management also supplies other metadata. This supplied metadata consists of IT template tables for the supported adapters, IT formulas, and IT report definitions. The installation process for SAS IT Resource Management stores the metadata in the Folders tree.

- The **SAS IT Resource Management** folder of the **Products** folder contains the metadata for the IT template tables of the supported adapters and the IT formulas. This metadata cannot be changed.
- The **SAS IT Resource Management** folder of the **Shared Data** folder contains the metadata for the 3.2 IT report definitions, IT data marts, and the IT formulas. This metadata can be changed, although changing the IT report definitions is not recommended. (If you need to change an IT report definition, copy it first and change the copy, not the supplied definition.)

CAUTION:

Do not make changes to the original IT report definitions in the Shared Data folder. For information about how to modify IT report definitions, see “Working with SAS Enterprise Guide Reports” in the *SAS IT Resource Management: Reporting Guide*.

- The **ITRM Transformations** folder of the Transformations tree contains the aggregation, gallery, information map, staging, and user-written staging transformations. These transformations can be dragged and dropped onto the process flow diagram of a job and further specified as needed.

Note: SAS IT Resource Management also uses folders, jobs, libraries, staged tables, aggregation table, and information maps. These objects are created by transformations that the user specifies (either directly or by means of the Adapter Setup wizard), the jobs that the user creates and runs, or by wizards (such as the New Folder, New IT Data Mart, or New Job wizards) that are invoked from the SAS IT Resource Management menu bar.

System Authentication When Logging On to the SAS Metadata Repository

To enable users to log on from SAS IT Resource Management to the SAS Metadata Server that contains the ITRM metadata, perform the following steps:

1. Invoke SAS IT Resource Management.
2. If you are already connected, from the menu bar, select **File** ⇒ **Connection Profile**. A message box displays a message that your application will be discontinued from its metadata server. It asks whether you want to continue. Click **Yes**.
3. The Connection Profile dialog box appears. In this dialog box, you can create a new connection profile or you can open an existing connection profile and modify it. After you specify the metadata profile that you want to use, click **OK**.
4. In the Login dialog box that opens, enter your user ID and password, and click **OK** to access the server that contains the metadata repository.

Note: If you previously chose to save the user ID and password for the particular connection profile, you are connected without being prompted to enter them.

When SAS IT Resource Management is installed, two groups of users are defined by default: IT Resource Management Administrators and IT Resource Management Users. Data administrators define additional users and groups as part of the setup tasks for SAS IT Resource Management. The metadata for all the logins for each user and group includes an authentication domain. Each user must have a login for the authentication domain that is associated with the relevant server definition. The user ID and password in that login must correspond to an account that has been established with the repository. In addition, this account must correspond to a valid account on the server. In other words, the user must already have an account on the target machine. Otherwise, the user will not be able to access any existing metadata. For more information about defining login metadata for users, see the *SAS Intelligence Platform: Security Administration Guide*. This document can be found in the Administration Documentation section of the SAS Intelligence Platform documentation that is available at: <http://support.sas.com/documentation/onlinedoc/intellplatform/index.html>.

Note: The permissions that govern access to the metadata can be set in the Access Control Template table in the **Authorization Manager** component of SAS Management Console. For more information, see the *SAS Intelligence Platform: System Administration Guide* and the *SAS Intelligence Platform: Security Administration Guide*.

Backing Up and Restoring the SAS Metadata Repository

For recommendations about backing up and refreshing the SAS Metadata Repository, see the *SAS Intelligence Platform: System Administration Guide*. When you back up the metadata server, it is important to also back up the data that is associated with the metadata objects that are contained in the repositories. The metadata cannot be used without the associated data. If you need to recover from a failure, the metadata that you restore must be synchronized correctly with the associated data.

Migrating from an Earlier Version of SAS IT Resource Management

If you worked with SAS IT Resource Management 3.1.1 or 3.2, you can migrate your metadata to the new environment. After you install and configure the current version of the software, you can migrate the metadata for all your IT data marts or for a single IT data mart. For information about these migration options, see *SAS IT Resource Management 9.3: Migration Documentation*. For more information about migrating the SAS system, see the *SAS Intelligence Platform: Migration Guide*.

If you created your own projects with SAS Enterprise Guide 4.1, 4.2, or 4.3, then you can use the Migration Wizard for SAS Enterprise Guide and the SAS Add-in for Microsoft Office to convert multiple projects at the same time to SAS Enterprise Guide 5.1. The Migration Wizard can change the paths in your SAS Enterprise Guide projects from the application server that was used with SAS 9.1.3 (SASMain) to the application server that is used with SAS 9.3 (SASITRM). If the paths were used in the SAS Enterprise Guide 4.1, 4.2, or 4.3 projects, the wizard can update the paths to the information maps. For information about the Migration Wizard, see “[Migrating SAS Enterprise Guide Projects to Work with SAS IT Resource Management 3.2 and Later](#)” on page 22.

Tools for Exploring SAS IT Resource Management Metadata

How to Access the SAS IT Resource Management Metadata

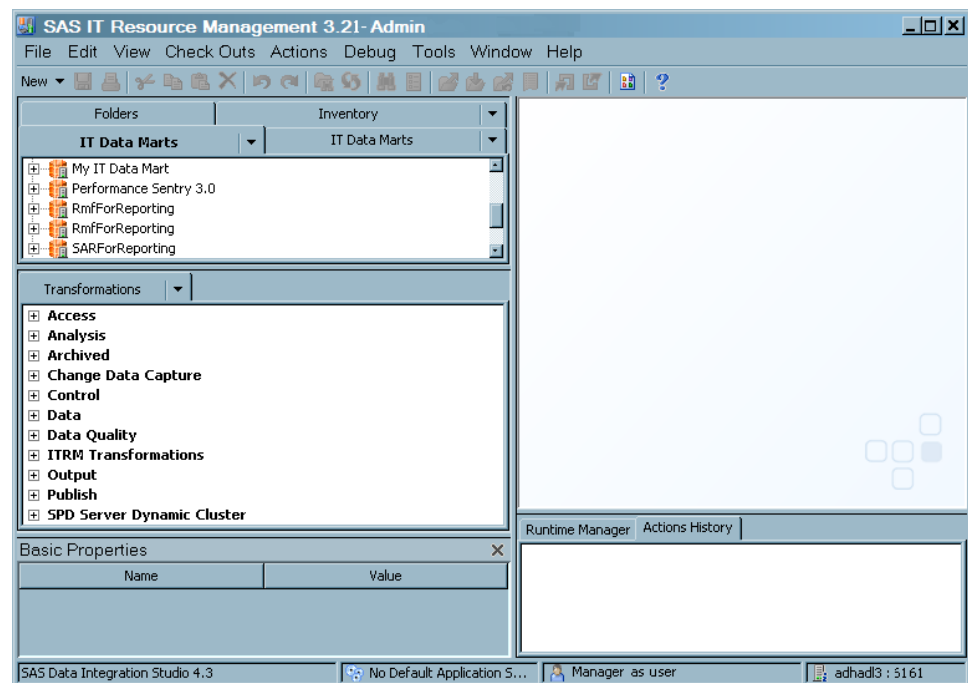
You can access the SAS IT Resource Management metadata in the SAS Metadata Repository by using the following methods:

- Explore the metadata repository from the SAS IT Resource Management Client. This method is the preferred way to access the SAS IT Resource Management metadata.
- Explore the metadata repository from the **Folders** tab of SAS Management Console.
- Browse the metadata repository from a Base SAS session. This method should be used only under the direction of SAS Technical Support.

Explore the Metadata Repository from the SAS IT Resource Management Client

The SAS IT Resource Management Client uses the data accessing and manipulation functions of SAS Data Integration Studio in order to populate its tables with data and to prepare that data for reporting.

Display 2.9 Main Window of the SAS IT Resource Management Client



The objects that SAS IT Resource Management uses are contained in folders in the Folders, IT Data Marts, and Transformations trees. For more detailed information about the many options and functions that are available within the SAS IT Resource Management client, see [“Features of the SAS IT Resource Management Client”](#) on page 36.

Browse the Repository from a Base SAS Session

Under the direction of SAS Technical Support, you can browse the repository from a Base SAS session. To do so, perform the following steps:

1. Enter *metabrowse* on the command line. The Metadata Server Configuration dialog box appears.

Display 2.10 Log on to the Metadata Server

2. Specify the appropriate logon information for the server that you want to access. Enter **Server Name**, **Port Number**, **Protocol**, **User Name**, and **Password**. Under Windows, the **User Name** can be domain qualified. For example, you might need to use *domainname\username*, where *domainname* is your domain name and *username* is your user ID.

Click **OK** to open the Metadata Browser.

3. To navigate through the metadata on this site, follow the directions of SAS Technical Support.

Metadata Server Memory Issues

Memory Issues with the SAS IT Resource Management Metadata Server

As the process size increases, the data administrator should delete from the SAS Metadata Repository any objects that are no longer needed or that have been marked for deletion by other processes. For best results, follow the recommendations for managing and maintaining your metadata server that are described in the *SAS Intelligence Platform: System Administration Guide*. This documentation is available at <http://support.sas.com/documentation/index.html> in the Administration Documentation section of that Web site.

Additional Requirements for z/OS

Using SAS IT Resource Management software on z/OS requires that batch jobs have a REGION of at least 1024M. A parallel ASSIZEMAX of 1024M is also required.

Chapter 3

Working with the SAS IT Resource Management Client

Overview of the SAS IT Resource Management Client	33
Introduction to the SAS IT Resource Management Client	33
Getting Started	34
Working with the Connection Profile Dialog Box	35
Using the Connection Profile Wizard	35
Features of the SAS IT Resource Management Client	36
About the Desktop	36
The Tree Views	38
The Job Editor	40
Properties of Jobs and Tables	41
Working with Wizards	43
Wizards Supplied by SAS IT Resource Management	43
Wizards Supplied by SAS Data Integration Studio	45
SAS IT Resource Management Transformations	45
Icons for SAS IT Resource Management Objects	46

Overview of the SAS IT Resource Management Client

Introduction to the SAS IT Resource Management Client

SAS IT Resource Management leverages and enhances SAS Data Integration Studio to administer and manage IT resource data. (The SAS IT Resource Management solution provides features and functionality in addition to that which is provided by SAS Data Integration Studio.) SAS Data Integration Studio is a visual design tool that provides a single point of control for managing data repositories and performing extraction, transformation, and loading (ETL) processes. It offers a variety of built-in transformations that enable you to perform the following tasks:

- build and manage data repositories
- construct and maintain data integration projects
- import data from and export data to a variety of sources
- create, modify, and deploy ETL jobs

The SAS IT Resource Management client uses SAS Data Integration Studio in order to create and manage metadata objects that define sources, targets, and the transformations that connect them. It also enables you to create the jobs that contain those transformations. SAS IT Resource Management uses metadata to create or retrieve SAS code that reads input from any data source that can be read by using SAS software. The metadata objects are stored in the SAS Metadata Repository that can be shared by multiple users.

SAS IT Resource Management enhancements to SAS Data Integration Studio provide additional object types, transformations, the Adapter Setup wizard, and other wizards that facilitate the process of generating and displaying analysis and report-ready data about your IT resources.

This chapter introduces the features and functions that are available with the SAS IT Resource Management client and provides basic information about the features and functions of SAS Data Integration Studio.

Note: To locate specific topics in the Help for SAS Data Integration Studio, you can use the Index or the Search feature.

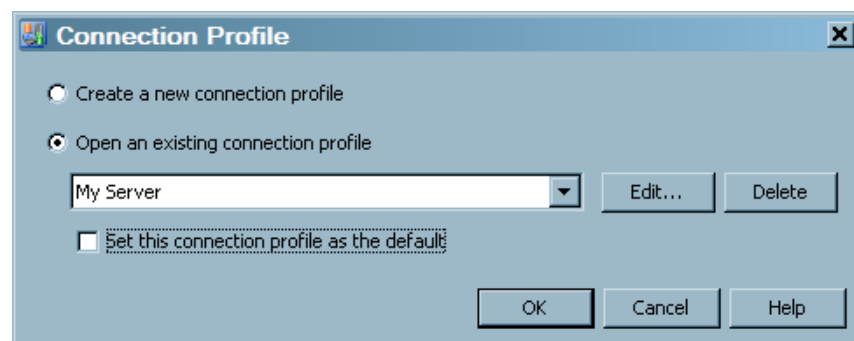
- From the menu bar of SAS IT Resource Management, select **Help** ⇒ **Contents and Index**.
- Click the **Index** tab or the **Search** tab. Then type the name of the topic in the **Find** field and press the ENTER key.

Getting Started

To access the online data manipulation and administrative services that are provided by the SAS IT Resource Management client, you must first log on to and connect to a metadata server:

1. To start SAS IT Resource Management from Microsoft Windows, select **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS IT Resource Management Client 3.2**. The following Dialog Box opens.

Display 3.1 Open a Connection Profile Dialog Box



2. Select (or create) a connection profile to connect to the server that contains the metadata for the SAS IT Resource Management objects such as jobs, templates, data, IT data marts, tables, information maps, report definitions, and transformations that you want to work with. The connection profile contains the information that identifies the server that contains that repository and the user who will be accessing the information from that repository.

Working with the Connection Profile Dialog Box

To manage your connection profiles, select **File** ⇒ **Connection Profile**. From the dialog box that opens, you can perform the following tasks:

- Create a new connection profile.
If you select this task and click **OK**, the Connection Profile wizard opens. Enter the information that is requested by the wizard. Click **Finish** to create the new connection profile.
- Open an existing connection profile.
If you select this task, from the drop-down list of existing connection profiles, select the profile that you want to open and click **OK**.
- Edit an existing connection profile.
If you click **Edit**, the Connection Profile wizard opens. Enter the changes to the information that is requested by the wizard. Click **Finish** to open the updated connection profile.
- Delete an existing connection profile.
If you click **Delete**, from the drop-down list of existing connection profiles, select the profile that you want to delete and click **Delete**. In the confirm deletion message box, click **Yes** to delete the connection profile.

Using the Connection Profile Wizard

The Connection Profile wizard guides you through the steps that are needed to build (or edit) a connection profile that enables you to connect to a metadata server. The wizard prompts you to enter (or revise) the following information:

- the name of the connection profile. If you want to set this profile as the default connection profile, check the corresponding box.
- the fully qualified name of the machine on which the repository server operates.
- the TCP port that the metadata server is listening to for connections. To enter a port number, type directly in the field.
- the user ID (and domain if applicable) that is used to log on to the metadata server. In order to specify the domain for this profile, you must also specify a value in the **Authentication Domain** field.

Note: You must authenticate to the operating environment where the metadata server is running.

- the password that is required for the specified user to log on to the metadata server. If you want to save the user ID and password in this connection profile, check the corresponding box.
- the authentication domain for the metadata server. If you specify a value for this field, the credentials that are specified for the metadata server are used to access servers or databases that are in the same authentication domain.

Note: You must authenticate to the operating environment where the metadata server is running.

- If you want to save the user ID and password in this profile, check the corresponding box.

- If you want to use the Integrated Windows authentication (single sign-on), check the corresponding box. When you check this box, the **Advanced** option is enabled.

Click the **Advanced** option to open the Advanced Settings dialog box. This dialog box enables you to specify the method for authenticating your credentials to the server and to SAS application servers. For information about how to specify the security package, the service principal name (SPN), and the security package list, click **Help**.

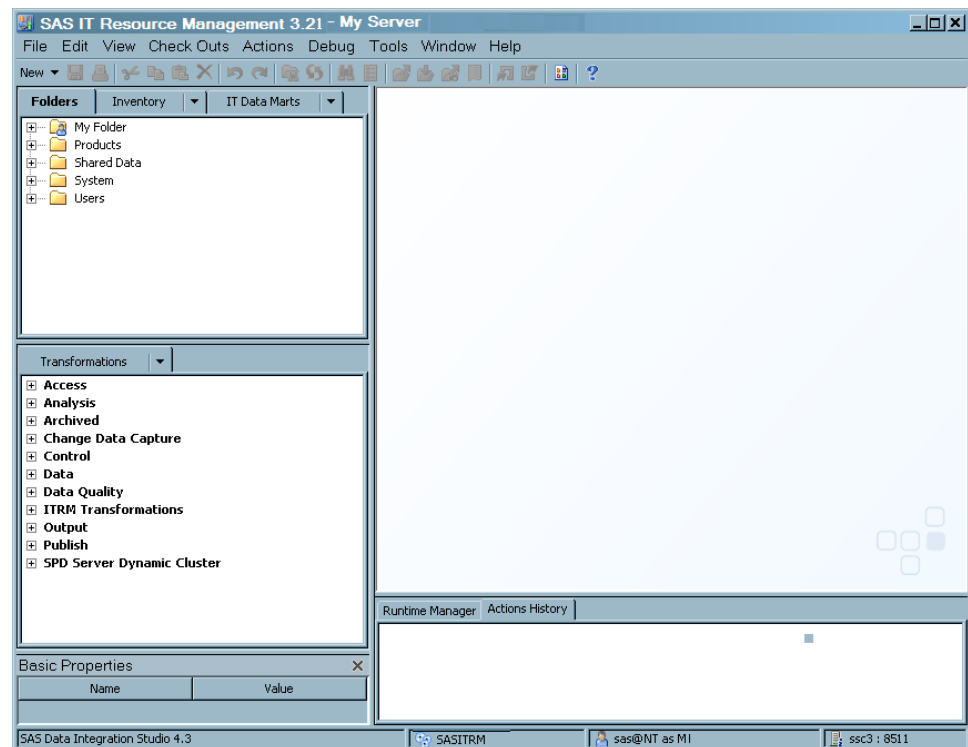
- After the required information has been entered, the wizard displays a summary page that lists the options that you specified for this connection profile. If necessary, you can go back to previous pages of the wizard and modify your specifications.
- When you click **Finish**, the connection profile is created (or updated).

Features of the SAS IT Resource Management Client

About the Desktop

After you open a connection profile, the SAS IT Resource Management desktop is displayed. (The SAS IT Resource Management client is powered by SAS Data Integration Studio.) The following display shows the desktop that opens by default.

Display 3.2 SAS IT Resource Management Desktop



The main components that are shown in the previous display are described in the following table.

Table 3.1 Desktop Components

Component	Location	Description
Title bar	Top of the desktop	Shows the current version of SAS IT Resource Management and the name of the current connection profile.
Menu bar	Under the title bar	Provides access to the drop-down menus. The list of active options varies according to the current work area and the type of object that you select. Inactive options are disabled or hidden.
Toolbar	Under the menu bar	Provides access to shortcuts for items on the menu bar. The list of active options varies according to the current work area and the type of object that you select. Inactive options are disabled or hidden.
Tree view	Left pane on the desktop	Provides access to the Folders, Inventory, Transformations, and IT Data Marts trees. It also provides access to the Basic Properties pane.
Basic Properties pane	Bottom of the left pane on the desktop	Displays basic properties of an object that is selected in the tree view. To display this pane, select View ⇒ Basic Properties from the desktop.
Status bar	Bottom of the desktop	<p>Displays the name of the currently selected object, the name of the default SAS Application Server if one has been selected, the login ID and metadata identity of the current user, and the name of the current SAS Metadata Server.</p> <p>To select a different SAS Application Server, double-click the name of that server to display a dialog box.</p> <p>If the name of the SAS Metadata Server turns red, the connection is broken. In that case, you can double-click the name of the metadata server to display a dialog box that enables you to reconnect.</p>
Runtime Manager	Under the Details pane	Displays the run-time status of the current job, the last time the job was executed in the current session, and the SAS Application Server that was used to execute the job. To display this pane, select View ⇒ Runtime Manager from the desktop.
Actions History	Under the Details pane	Displays low-priority errors and warnings. To display this pane, select View ⇒ Actions History from the desktop.

The Tree Views

About Trees

Trees facilitate how you work with SAS IT Resource Management. They provide you with easy access to the folders, transformations, and metadata objects that are registered with the SAS Metadata Server. In addition to the Folders, Inventory, and Transformations trees that are provided by SAS Data Integration Studio, the SAS IT Resource Management enhancements to SAS Data Integration Studio provide the IT Data Marts tree.

You can change the way the trees are displayed in the left panel of the desktop. To do so, select one of the options from the drop-down list. These options enable you to move the tree up, down, to the right, or to the left in the left panel.

IT Data Marts Tree

The IT Data Marts tree displays the IT data marts that can be accessed by the current connection profile. IT data marts contain the metadata for the jobs, libraries, tables, information maps, and other elements that support the ETL processes and reporting of IT resource data. As such, the IT Data Marts tree is the primary starting point for most of the SAS IT Resource management processes. For best results, perform your IT resource management tasks from within this tree.

For more information about IT data marts, see [“About the IT Data Mart” on page 51](#).

Folders Tree

The Folders tree organizes metadata into folders that are shared across a number of SAS applications. **My Folder** and **Shared Data** are the folders that you use most of the time. The following folders are supplied by default:

- **My Folder** is the private folder of the user who is currently logged on. It can be used to store metadata that is not available to other users. Metadata in this folder can be viewed only by the user who is currently logged on and by users who have unrestricted access.
- **Products** is the folder that contains the software that is installed with SAS Data Integration Studio and SAS IT Resource Management. Along with other objects, this folder contains the template tables for the adapters that are installed and the supplied formulas that can be used to compute new columns from existing columns of data. The contents of this folder cannot be modified. However, they can be viewed and copied to another folder.
- **Shared Data** is the folder that contains data that can be accessed by any user who is logged on. It contains the supplied SAS Enterprise Guide report definitions and the formulas that are installed with SAS IT Resource Management. This folder can also contain IT data marts that are intended to be accessed by multiple users.
- **System** is the folder that contains applications, channels, subscribers, services, and the dialog boxes of the types of objects that are installed and used by SAS Data Integration Studio and SAS IT Resource Management.
- **Users** is the folder that contains the private folders of users who are connected to the active metadata server.

Inventory Tree

The Inventory tree displays metadata for objects that are registered on the current metadata server, such as tables and libraries. Metadata in this tree can be stored in folders that group metadata by type, such as Table, Library, Job, Information Map, and so on.

Transformations Tree

The Transformation tree displays transformations that can be dragged and dropped onto the process flow diagrams of SAS Data Integration Studio jobs. SAS IT Resource Management provides the **ITRM Transformations** folder that contains the Aggregation and Information Map transformations. In addition, this folder contains the staging transformations for each of the supplied adapters.

The transformation's Properties dialog box enables you to view or update the metadata for a transformation in a SAS Data Integration Studio job. One way to display this dialog box is to open a job on the **Diagram** tab of the Job Editor, right-click a transformation on the **Diagram** tab, and click **Properties** in the pop-up menu. The Properties dialog box for most transformations has one or more tabs that are unique to that transformation. The following table describes the purpose of the common tabs for a transformation. For more information about each tab, see the Help for that tab.

Table 3.2 Common Tabs in a Transformation Properties Dialog Box

Tab	Description
General	Enables you to enter general information that identifies and describes the transformation.
Mappings	Enables you to review and modify the mappings for the transformation. <i>Note:</i> This tab is not used by SAS IT Resource Management transformations.
Options	Enables you to review and modify options for the transformation.
Table Options	Enables you to review and modify table options for the transformation.
Code	Enables you to review and modify the code that is generated for the transformation.
Precode and Postcode	Enables you review and modify user-written code that is inserted at the beginning or end of the transformation.
Parameters	Enables you to review and modify parameters for the transformation.
Notes	Enables you to review and modify notes for the transformation.
Extended Attributes	Enables you to review and modify extended attributes for the transformation.

Basic Properties Pane

The Basic Properties pane displays the basic properties of an object that is selected in a tree view. The Basic Properties pane is an optional pane that can be displayed on the left side of the desktop. To display or hide this pane, select or deselect **View** ⇒ **Basic Properties** from the menu bar.

Checkouts Tree

The Checkouts Tree displays metadata that has been checked out for update, as well as any new metadata that has not been checked in. The metadata is displayed automatically when you are working under the change management function of SAS Data Integration Studio.

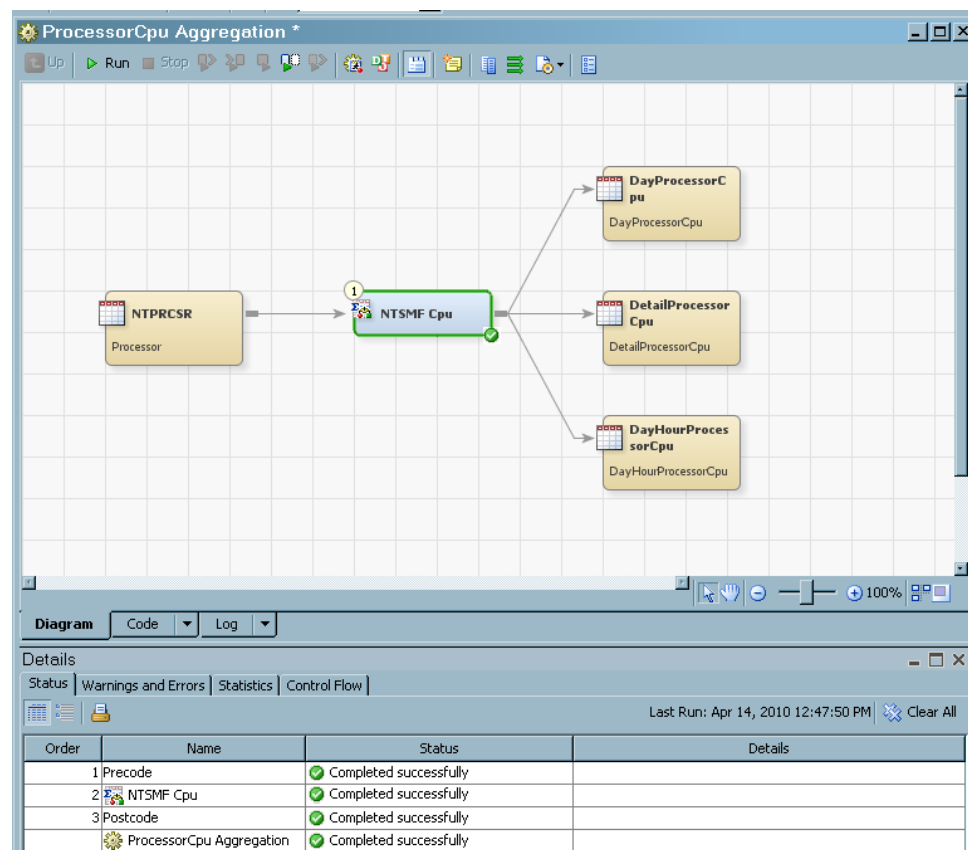
Note: This tree is not used by SAS IT Resource Management.

The Job Editor

Job Editor Window

The Job Editor window enables you to create, maintain, and troubleshoot SAS IT Resource Management jobs. To display this window, right-click a job in the tree view and select **Open**. The following display shows a sample Job Editor window.

Display 3.3 Sample Job Editor Window



To open an existing job in the Job Editor window, navigate to the folder where it is located and double-click it. To create a new job, from the menu bar select **New** ⇒ **Job**.

In order to maintain unique job names within the repository, the name of the new job is a concatenation of "New Job" and a five-digit number (for example, "New Job 28424".)

Note: You can rename a job to a more meaningful name. For information about how to rename a job, see [“Rename a Job” on page 273](#).

The job opens in the Job Editor window. In the Job Editor window, you can add, delete, and modify the transformations and tables of a job.

Job Editor Tabs

The following table describes the main tabs in the Job Editor window.

Table 3.3 Job Editor Tabs

Tab	How to Display the Tab	Description
Diagram	Always displayed.	Used to build and update the process flow for a job.
Code	Select Tools ⇒ Options from the desktop. On the General tab, select Show Code Tab .	Used to review or update code for a job.
Log	Select Tools ⇒ Options from the desktop. On the General tab, select Show Log Tab .	Used to review the log for a submitted job.
Output	Select Tools ⇒ Options from the desktop. On the General tab, select Show Output Tab .	Used to review the output of a submitted job.

Use the Details pane to monitor and debug a job in the Job Editor. To open this pane, select **View** ⇒ **Details** from the desktop.

Use the Runtime Manager pane to display the run-time status of the current job, the last time the job was executed in the current session, and the SAS Application Server that was used to execute the job. This information is available as long as the job is active. To open this pane, select **View** ⇒ **Runtime Manager** from the desktop.

Properties of Jobs and Tables

Job Properties

The job Properties dialog box enables you to view or update the metadata for a SAS Data Integration Studio job. One way to display this dialog box is to right-click a job in the IT Data Marts tree, Folders tree, or Inventory tree, and click **Properties** in the menu that is displayed. The following table describes the purpose of each tab in a job's Properties dialog box. For more information about each tab, see the Help for that tab.

Table 3.4 *Tabs in a Job Properties Dialog Box*

Tab	Description
General	Enables you to enter general information that identifies, describes, and locates the job.
Code	Enables you to review and modify the code that is generated for the job.
Precode and Postcode	Enables you to review and modify user-written code that is inserted at the beginning or end of the job.
Status Handling	Enables you to review and modify status handling conditions and actions for the job. (Some transformations have this tab as well.)
Parameters	Enables you to review and modify parameters for the job.
Options	Enables you to review and modify options for the job.
Notes	Enables you to review and modify notes for the job.
Extended Attributes	Enables you to review and modify extended attributes for the job.
Authorization	Enables you to review and modify metadata access settings for the job.

Table Properties

The table's Properties dialog box enables you to view or update the metadata for the table. One way to display this window is to right-click a table in the Folders tree or the Job Editor and click **Properties**. The next table describes the purpose of each tab in a table's Properties dialog box. For more information about each tab, see the Help for that tab.

Table 3.5 *Tabs in a Table Properties Dialog Box*

Tab	Description
General	Enables you to enter general information that identifies and describes the table.
Columns	Enables you to maintain column metadata.
Indexes	Enables you to review, add, and modify indexes on table columns.
Keys	Enables you to review, add, and modify key columns.
Parameters	Enables you to review and modify parameters for the table.

Tab	Description
Physical Storage	Enables you to specify the format and location of a table.
Options	Enables you to review and modify options for the table.
Notes	Enables you to review and modify notes for the table.
Extended Attributes	Enables you to review and modify extended attributes for the table.
Authorization	Enables you to review and modify metadata access settings for the table.

Working with Wizards

Wizards Supplied by SAS IT Resource Management

In addition to the wizards that are supplied by SAS Data Integration Studio, SAS IT Resource Management provides the following wizards:

The following wizards are invoked from the **New** menu:

- defining ETL jobs

The Adapter Setup wizard helps you define the ETL jobs that read, stage, aggregate, map, and generate information maps and reports for the IT resource data that an adapter loads. This wizard also helps you define the report jobs that generate gallery reports.

The Adapter Setup wizard produces data acquisition, staging, aggregation, information map, and report jobs. Configuration options of minimal, typical, and full are available to define the quantity of jobs generated by the Adapter Setup wizard and are good choices for most user's needs. You can also specify that only the staging job is generated. However, your site might have other requirements that limit or expand the number and type of jobs that should be generated. Therefore, you should examine the results of the Adapter Setup wizard and possibly add or delete elements of the jobs, transformations, or tables in order to accommodate the needs of your site or organization.

To invoke the Adapter Setup wizard from the SAS Data Integration Studio toolbar, select **New** ⇒ **IT Resource Management** ⇒ **Adapter Setup**.

For more information about the Adapter Setup wizard, see [“About the Adapter Setup Wizard” on page 177](#).

- New Formula wizard

SAS IT Resource Management provides standard formulas that can be used to generate a computed column. To locate these formulas, navigate to the **Shared Data** folder in the Folders tree and select **SAS IT Resource Management** ⇒ **IT**

Formulas. You can copy these formulas, modify them, or create new customized formulas by using the New Formula wizard.

To invoke the New Formula wizard, navigate to the IT Data Marts tree. Then select **New** ⇒ **IT Resource Management** ⇒ **Formula**.

For more information about formulas, see [“About Formulas” on page 97](#).

- New IT Data Mart wizard

The New IT Data Mart wizard enables you to create a new IT data mart, which is used to contain the metadata for the jobs, libraries, tables, information maps, and other elements that support the ETL processes of and reporting on IT performance data.

To invoke the **New IT Data Mart** wizard, navigate to the IT Data Marts tree. Then select **New** ⇒ **IT Resource Management** ⇒ **IT Data Mart**.

Note: You can also invoke the New IT Data Mart wizard from the first page of the Adapter Setup wizard by selecting the **New IT data mart** option.

For more information about IT data marts, see [“About the IT Data Mart” on page 51](#).

The following wizards are supplied by SAS IT Resource Management, but are not invoked directly from the **New** menu:

- Aggregation wizards

Aggregation tables contain IT resource data that has been classified, summarized, or aged according to the specifications of an Aggregation transformation.

To invoke the aggregation table wizards, open (or create) a job that contains an Aggregation transformation. Right-click the Aggregation transformation and select the **Add Aggregation Table** option. You can then choose one of the following options:

- Summarized Aggregation Table, which enables you to directly specify the variables (and their associated properties) that are to be aggregated.
- Simple Aggregation Table, which is not summarized. The new data is simply appended to the aggregation table.
- Aggregations through Template, which creates aggregation tables by using the provided aggregation templates.

For more information about aggregating data, see [“Working with Aggregation Transformations” on page 170](#).

- Maintain Staged Tables

The Maintain Staged Tables wizard enables adapter updates that are delivered by SAS to be applied to your IT data marts. It enables you to update the staged tables of an IT data mart, based on the revisions that have been made in the IT resource collection software and files.

To invoke the Maintain Staged Tables wizard, navigate to the IT Data Marts tree. Then right-click the IT data mart whose tables are to be updated, and select the **Maintain staged table** option.

For more information about IT data marts and maintaining their associated tables, see [“About the IT Data Mart” on page 51](#).

- Staged Table wizard

A staged table is a SAS IT Resource Management table that contains data that has been extracted from an input data store and rendered into a form that is suitable for further transformation. It can be created either by the Adapter Setup wizard or by the Staged Table wizard.

To invoke the Staged Table wizard, open (or create) a job that contains a staging transformation. Right-click the staging transformation and select the **Add Staged Tables** option.

For more information about staged tables, see [“Working with Staging Transformations” on page 107](#).

Wizards Supplied by SAS Data Integration Studio

After you are connected to the server that you want to use, you can access the wizards that SAS Data Integration Studio provides. These wizards can create the new objects that can help you work with your IT performance data. From the menu on the SAS IT Resource Management toolbar select **New** to open the wizards, dialog boxes, and windows that create the following objects:

- Folder
- Job
- Table
- Transformation
- External File
- Library
- Document
- Note
- IT Resource Management

Note: This option enables you to access the wizards that the SAS IT Resource Management enhancements to SAS Data Integration Studio provide.

- Cube
- OLAP Schema

Note: Help is available on the pages of these wizards, dialog boxes, and windows.

SAS IT Resource Management Transformations

SAS IT Resource Management supplies many transformations that enable you to work with your IT data. To specify a transformation, you must first open (or create) a job. This action creates a process flow diagram on the **Diagram** tab of the Job Editor window. Transformations can be dragged and dropped onto the process flow diagram of the job, and their properties, data sources, and resulting tables can then be specified according to your requirements. The Adapter Setup wizard specifies the necessary transformations programmatically, or you can do so by working directly with the transformations.

In addition to the standard SAS Data Integration Studio transformations that are accessible from the Transformations tree, SAS IT Resource Management provides the following transformations:

Type of Transformation	Description
Aggregation	Specifies how a staged (or other) table is to be aggregated into a simple or summarized aggregation table.
Data Acquisition	Specifies how to extract a subset of VMware data from a database.
Gallery	Specifies the gallery reports that are to be generated from information maps or aggregation tables and how these reports are to be stored on the SAS Content Server.
Information Map	Creates information maps from aggregation tables.
Staging (one for each supported adapter)	Specifies how raw data is to be processed and loaded into staged tables for the supported adapters.
User-Written Staging	Specifies how raw data is to be processed and loaded into staged tables for data that is not supported by IT Resource Management adapters.











To locate these transformations, select the Transformations tree and navigate to the **ITRM Transformations** folder. For information about the properties of transformations, see [“Transformations Tree” on page 39](#).

Icons for SAS IT Resource Management Objects

The following table lists icons that are used by objects in SAS IT Resource Management.

Table 3.6 SAS IT Resource Management Icons

Object Name	Icon
Adapter transformation	
Aggregation transformation	
Column (Character contents)	
Column (Numeric contents)	
Deployed Job	
External file	
Folder	

Object Name	Icon
IT Formula	
Gallery transformation	
Information Map	
IT Data Mart	
Job	
Library	
IT Report Definition	
IT Resource Management Desktop	
Table	
IT Template	

Part 2

Accessing and Processing IT Data

<i>Chapter 4</i> IT Data Mart	51
<i>Chapter 5</i> Adapters	67
<i>Chapter 6</i> Formulas	97
<i>Chapter 7</i> Staging the Data	105
<i>Chapter 8</i> Aggregating the Data	131
<i>Chapter 9</i> Adapter Setup Wizard	177
<i>Chapter 10</i> Information Maps	207
<i>Chapter 11</i> User-Written Staging Code	233
<i>Chapter 12</i> Jobs	263

Chapter 4

IT Data Mart

About the IT Data Mart	51
What Is an IT Data Mart?	51
How Is an IT Data Mart Created?	52
About the IT Data Mart Wizard	52
Where Can the IT Data Mart Be Stored?	52
How Can the IT Data Mart Be Accessed?	52
How Is the IT Data Mart Populated?	52
What Are the Contents of an IT Data Mart?	53
Working with IT Data Marts	56
Introduction to Working with IT Data Marts	56
Copy and Paste an IT Data Mart	56
Create an IT Data Mart	57
Delete an IT Data Mart	58
Erase an IT Data Mart	58
Modify an IT Data Mart	59
Purge the Tables of an IT Data Mart	59
Rename an IT Data Mart	60
Importing, Exporting, and Promoting Metadata	60
About Importing, Exporting, and Promoting Metadata	60
About Promoting an IT Data Mart	61
How to Export an IT Data Mart	61
How to Import an IT Data Mart	62

About the IT Data Mart

What Is an IT Data Mart?

An IT data mart is a key component of SAS IT Resource Management. It is the logical collection of the jobs, data, information maps, tables, and other elements that support the processing and analysis of IT data. IT data marts are set up, managed, and administered by the data administrator using the wizards and user interfaces that are available in SAS IT Resource Management. Multiple IT data marts can be created in order to help organize a site's data. For example, a site's data administrator might set up a separate data mart for each of the operating system types or business areas that exist within the corporate enterprise.

How Is an IT Data Mart Created?

An IT data mart is created by invoking the New IT Data Mart wizard. This wizard can be invoked in two ways:

- from the SAS IT Resource Management menu bar.

From the SAS IT Resource Management menu bar, select **New** ⇒ **IT Resource Management** ⇒ **IT Data Mart** to open the New IT Data Mart wizard.

- from the Adapter Setup wizard.

From the SAS IT Resource Management menu bar, select **New** ⇒ **IT Resource Management** ⇒ **Adapter Setup** to open the Adapter Setup wizard. On the first page of that wizard, you can click **New IT data mart** to open the New IT Data Mart wizard. (You can also select an existing IT data mart from this page.)

About the IT Data Mart Wizard

The New IT Data Mart wizard guides you through the steps to create an IT data mart. You are prompted to specify the name and description of the IT data mart, the default application server, and the default root path. You can also change the default location where the IT data mart is to be stored.

On completion, the New IT Data Mart wizard generates the main root folder for the IT data mart, the IT data mart itself, an **Administrative** subfolder, and the administrative library, **ADMIN**.

Note: The name of the Admin library is programmatically generated by the concatenation of **Admin** and a random number, **nnnn**.

Where Can the IT Data Mart Be Stored?

When you create an IT data mart, you can specify a location where it will be stored.

The **Location** field on the first page of the New Object wizard enables you to specify where your IT data mart is to be stored. By default, your IT data mart is stored in the **SAS IT Resource Management/IT Data Marts** subfolder of the **Shared Data** subfolder. However, you can change this location to any location to which you have Write privileges.

How Can the IT Data Mart Be Accessed?

You can access an IT data mart from the subfolders that are listed on the IT Data Marts tree, the Folders tree, or the Inventory tree of SAS IT Resource Management. To work most effectively with your IT data mart, access it from the IT Data Marts tree.

How Is the IT Data Mart Populated?

IT data marts can be populated in two ways.

- The standard way to populate the IT data mart is through the Adapter Setup wizard. For information about the Adapter Setup wizard, see [“Adapter Setup Wizard” on page 177](#).

- You can also populate an IT data mart with jobs, tables, libraries, and other objects by explicitly adding individual objects, by using the menu actions and menu bar actions.

What Are the Contents of an IT Data Mart?

IT data marts can contain the following objects:

- Admin library

The Admin library is stored in the **Administrative** folder of the IT data mart. You can store any content, including SAS formats, in this library. The Admin library also stores adapter resource lookup tables for the staging transformations of several adapters, such as the SNMP or IBM SMF adapters.

Note: If the IT data mart is created outside of the Adapter Setup wizard (that is, by invoking the IT Data Mart wizard directly from the **New** menu of SAS IT Resource Management), the **Administrative** folder with its ADMIN library, is the only subfolder that is created.

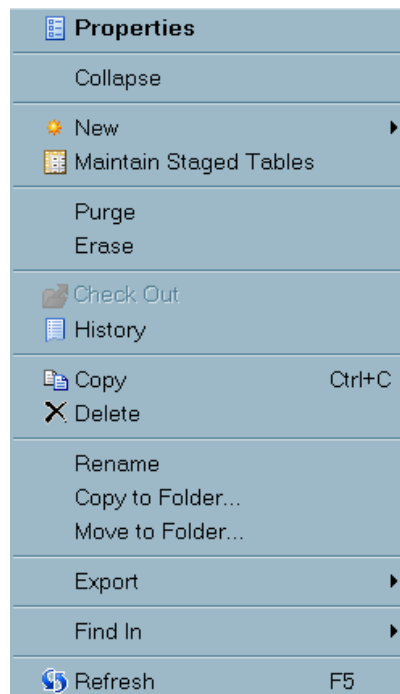
CAUTION:

Do not delete this library. Deleting this library can cause processing errors.

- IT data mart object (🏠)

The IT data mart object enables you to perform certain tasks, which are available when you right-click the object.

Display 4.1 Actions Available by Right-clicking an IT Data Mart



In the IT Data Marts tree, the IT data mart object and its root folder object are the same; in the Folders tree, the IT data mart object, and its root folder object appear as two separate objects. Use the IT Data Marts tree for most IT Resource Management tasks that involve IT data marts.

CAUTION:

For best results, use the IT Data Marts tree for most IT Resource Management tasks that involve IT data marts. Working in the IT Data Marts tree ensures that the metadata is updated properly.

- adapter-based subfolders

A separate subfolder is generated each time you successfully execute the Adapter Setup wizard. The name of the subfolder is based on the adapter that you selected during the Adapter Setup wizard. Each adapter-based subfolder contains a **Domain Categories** subfolder and a **Staging** subfolder.

- Separate subfolders in the **Domain Categories** subfolders contain the objects such as jobs and tables for each domain category that you specified. The name that is assigned to this subfolder is the name of the domain category that you selected. The objects that are added for a domain category can include the output aggregation tables, libraries, information maps, and the aggregation, information map, and reporting jobs that are associated with that domain category.

Note: The physical tables are generated when their corresponding jobs are run. These tables do not have to be stored in a fixed location with respect to the root location of the IT data mart. You can store these tables in any location to which you have Write privileges.

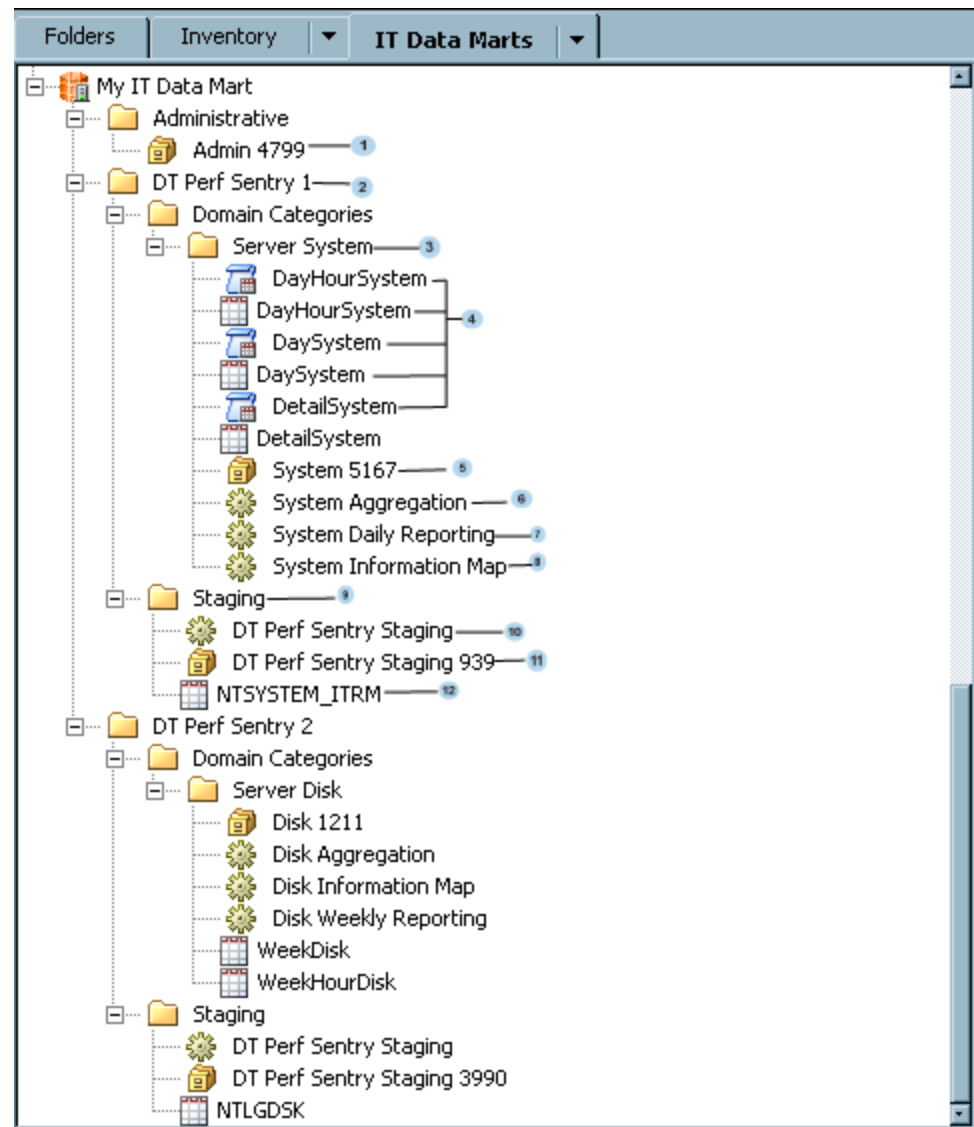
Note: The metadata for information maps is not generated until the information map job is run. The metadata for the information maps does not have to be stored in a fixed location with respect to the root location of the IT data mart. You can store this metadata in any location to which you have Write privileges.

- A **Staging** subfolder is generated for each instance of the adapter that you selected. This subfolder contains the staging job, library, and staged tables for that adapter. (For some adapters, a spin library is also generated.)

The name that is assigned to the adapter-based subfolder is generated based on the name of the selected adapter.

The following figure displays the IT Data Mart tree view of My IT Data Mart.

Figure 4.1 Example of an IT Data Mart



- 1 Admin library that was generated for MY IT Data Mart
- 2 Folder that is generated for the first instance of running the Adapter Setup wizard for the DT Perf Sentry adapter
- 3 Folder that contains objects related to the Server System domain category for DT Perf Sentry
- 4 Tables and information maps for the aggregations that are connected with the Server System domain
- 5 Library that contains all aggregation tables that are connected with the Server System domain
- 6 Job that generates the aggregation tables that are connected with the Server System domain
- 7 Job that generates reports that are connected with the Server System domain

- 8 Job that generates information maps that are connected with the Server System domain
- 9 Folder that contains the staging library, tables, and job that are connected with DT Perf Sentry
- 10 Staging job for DT Perf Sentry
- 11 Library that contains all staged tables that are populated by the DT Perf Sentry Staging job
- 12 Staged table that is populated by the DT Perf Sentry Staging job

In the preceding figure, the IT data mart called MY IT Data Mart also contains metadata for two instances of the DT Perf Sentry adapter:

- DT Perf Sentry 1, the first instance of using the Adapter Setup wizard for DT Perf Sentry
- DT Perf Sentry 2, the second instance of using the Adapter Setup wizard for DT Perf Sentry

Folders, libraries, jobs, and staging information are generated for these, and any other adapters, according to the [Figure 4.1 on page 55](#). However, the IT data mart is populated with the information maps when the information map job runs.

Note: Some adapters require a SPIN library that manages transactions or events that are not yet complete. If required, SAS IT Resource Management creates and registers the library.

Working with IT Data Marts

Introduction to Working with IT Data Marts

You can create, delete, erase, modify, purge, and rename an IT data mart. For best results, perform these functions from the IT Data Marts view.

Copy and Paste an IT Data Mart

To copy an IT data mart and then paste it into another location, perform the following steps:

1. Determine the folder into which you want to paste the IT data mart. If you want to create a new folder for the IT data mart, from the Folders tree, click **New**. From the drop-down list, select **Folder**. The New Folder wizard opens. You can retain the default name of the new folder, or you can enter a new name. You can also retain the default location, or you can click **Browse** to navigate to another location. Select the location that you want to use. Click **OK** to create the new folder.
2. Copy the IT data mart.
To do so, from either the Folders tree or the IT Data Marts tree, right-click the IT data mart that you want to copy and select **Copy**.
3. Paste the IT data mart into its new location.

To do so, in the Folders tree, right-click the folder into which you want to paste the IT data mart. Then select one of the following actions:

- Select **Paste** to paste the IT data mart into the selected folder. (The Paste action pastes only the IT data mart itself and a limited number of associated objects, such as the Administrative folder.)
 - If you want to create a new IT data mart that contains other components of the IT data mart that you copied, select **Paste Special**. The Paste Special window appears. In this window, you can specify which contents of the IT data mart you want to paste into the new IT data mart.
4. Click **Finish** to create the new IT data mart.

Create an IT Data Mart

To create an IT data mart, perform the following steps:

1. From the SAS IT Resource Management menu bar, select **New** ⇒ **IT Resource Management**. From the list that is displayed, select **IT Data Mart**. The New IT Data Mart wizard opens.
2. On the General information page of the wizard, specify the name, description, and location of the IT data mart.
 - The name of the new data mart is required. The **Name** field can contain a maximum of 60 characters. The wizard does not permit you to enter invalid characters. The name must be unique within its parent folder. For best results, the name should be unique within the metadata server.
 - A value is optional in the **Description** field. The **Description** field can contain a maximum of 200 characters.
 - The **Location** field contains the default location for the IT data mart. A value is required in this field. This location is typically in the **Shared Data\SAS IT Resource Management** folder of the repository where you are logged on. You can modify this field to specify any location to which you have Write privileges. The wizard does not permit you to enter invalid characters.

Note: Make sure that the folder where the IT data mart is to be stored is not the root folder, that is in the **/Users/<userID>** folder. If you want to store the IT data mart under the **/Users/** folder, make sure that you specify the folder location as **/Users/<userID>/MyFolders**.

3. On the Application server information page, specify the default application server.

Application servers store and execute SAS code. The default application server that is specified with this wizard accesses metadata, DBMS tables, and external files on remote hosts so that it can build and support an IT data mart. A value is required in this field.

Select the default application server from the drop-down list of currently defined SAS Application servers. Then click **Test Connection** to make sure that you are connected to the server that you specified. If you are not connected at this time, you are prompted to enter your user ID and password in order to access the server.

Note: Application servers are defined when the SAS IT Resource Management software is installed and configured. Use SAS Management Console to modify this list of available servers.

4. On the Root path information page, specify the path that will be used for the creation of the physical libraries. A value is required in this field. The path that you specify should be a fully qualified path - relative paths are not supported.

You can type the path in the field, or you can click **Browse** and select a root path from the Select a Directory dialog box that displays the directory structure on your file system. IT data marts with invalid root path will not be created and an error message that explains the problem will be displayed. This physical path is not validated. Make sure that each IT data mart uses a different path. The value for this field cannot contain the single quotation character (').

Browse is disabled when the selected application server is running on a machine that uses the z/OS operating system. When entering a z/OS path manually, you can use either a prefix for a traditional z/OS filesystem path (**MY.DATAMART**) or a root directory in the UFS hierarchical file system (**/u/myname/datamart**).

Note: The documentation for SAS on z/OS has traditionally used the terms UNIX System Services (USS) and hierarchical file system (HFS) to refer to the UNIX file system on z/OS. The SAS 9.3 documentation uses the terms UNIX file system and UFS to refer to this file system. In addition to the original HFS implementation, the z/OS operating system also provides another UNIX file system known as the z/OS file system (zFS). zFS, which provides certain performance and manageability benefits, is functionally equivalent to HFS from the perspective of a SAS user.

5. On the Summary page, the wizard displays the specifications that will be used to create the new IT data mart. If they are satisfactory, click **Finish** to create the new IT data mart.

Delete an IT Data Mart

Deleting an IT data mart removes all the metadata that is associated with the IT data mart. It does not remove the physical tables. To delete an IT data mart and all its contents, perform the following steps:

1. On the IT Data Marts tree of SAS IT Resource Management, right-click the IT data mart that you want to delete. (Make sure that the jobs that are located in that IT data mart are not open in a **Diagram** tab of the Job Editor window.)

Note: You can also access the Delete function from the Folders tree of SAS IT Resource Management. However, performing tasks on the IT data mart from any tree other than the IT Data Marts tree is not recommended because the metadata that is associated with the IT data mart might not be properly updated.

2. Click **Delete**. A message box appears that asks you to confirm that you want to delete the objects from the selected IT data marts and all their contents.
3. To continue with the deletion, click **Yes**. To cancel the deletion, click **No**.

CAUTION:

Do not delete an IT data mart from the Folders tree. If you want to delete an IT data mart, always do so from the IT Data Marts tree.

Erase an IT Data Mart

Erasing an IT data mart deletes all the SAS libraries and their contents from the metadata. It also deletes all empty directories at the root level of the IT data mart. (The

Erase action does not delete any directories that have contents.) It also deletes all the metadata that is in the IT data mart. To erase an IT data mart and all its contents, perform the following steps:

1. On the IT Data Marts tree of SAS IT Resource Management, right-click the IT data mart that you want to delete. (Make sure that the jobs that are located in that IT data mart are not open in a **Diagram** tab of the Job Editor window.)

Note: You can also access the Delete function from the Folders tree of SAS IT Resource Management. However, performing tasks on the IT data mart from any tree other than the IT Data Marts tree is not recommended because the metadata that is associated with the IT data mart might not be properly updated.

2. Click **Erase**. A message box appears that asks you to confirm that you want to erase the objects from the selected IT data marts and all their contents.
3. To continue with the deletion, click **Yes**. To cancel the erase action, click **No**.

Modify an IT Data Mart

You can modify some of the properties of an IT data mart.

To do so, perform the following steps:

1. On the IT Data Marts tree of SAS IT Resource Management, right-click the IT data mart that you want to modify. Select **Properties**. (Make sure that the jobs that are located in that IT data mart are not open in a **Diagram** tab of the Job Editor window.)

2. • Select the **General** tab to view and modify the **Name**, **Description**, and **Location** fields.

Note: If you rename an IT data mart to the same name as an existing IT data mart within the same application server, the following message is displayed: “IT Data Mart folder with this name already exists in the selected folder location. Change the name or select a different folder location.”

- Select the **Extended Attributes** tab to view and modify the fields on this tab.
Use the **Add** and **Delete** buttons to add and delete fields. In general, you should not change or delete any supplied fields on this tab.
- Select the **Advanced** tab to view metadata for the IT data mart. In general, modifying these fields can cause processing errors. However, you can modify the **Name** and **Description** fields.

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

Note: If you change anything about a job, you should redeploy all the jobs that are contained in the IT data mart. These changes can include adding or deleting tables or columns, changing the locations of tables, and changing a parameter on a transformation.

Purge the Tables of an IT Data Mart

Purging an IT data mart removes the contents of the tables in that IT data mart. It does not remove the metadata or the physical tables. To purge the tables of an IT data mart, perform the following steps:

1. On the IT Data Marts tree of SAS IT Resource Management, right-click the IT data mart whose table you want to purge.
2. From the drop-down menu that displays, click **Purge**.
3. Review the list of tables that are displayed in the confirmation dialog box. Click **Yes** to delete the tables. Otherwise, click **No**.

Rename an IT Data Mart

Note: If the renaming of the IT data mart includes moving it to another folder, perform the **Move to Folder** task in the Folders view. Then, perform the renaming task in the IT Data Marts view.

To rename an IT data mart, perform the following steps:

1. On the IT Data Marts tree of SAS IT Resource Management, right-click the IT data mart that you want to rename. (Make sure that the jobs that are located in that IT data mart are not open in the **Diagram** tab of the Job Editor window.)
2. From the drop-down menu, click **Rename**. The name of the IT data mart is highlighted.
3. Enter the new name of IT data mart in the highlighted area.

Note: If you rename an IT data mart to the same name as an existing IT data mart within the same application server, the following message is displayed: “IT Data Mart folder with this name already exists in the selected folder location. Change the name or select a different folder location.”

Alternatively, you can rename the IT data mart by right-clicking the IT data mart and selecting **Properties**. Select the **General** tab. Enter the new name of the IT data mart in the **Name** field. If you rename an IT data mart to the same name as an existing IT data mart, a message is displayed that warns that the new name is already in use for another IT data mart.

Note: In any job, if you have explicitly specified the path to an IT data mart, and you subsequently rename that IT data mart, then you should change the name of the IT data mart in that explicitly specified path. In addition, if you change any other folder name that is used in an explicit path reference, you should ensure that the explicitly specified path uses the changed folder name.

CAUTION:

Redeploy the jobs that are contained in the IT data mart. Redeploying jobs is necessary because renaming an IT data mart changes the metadata path of anything that is associated with that IT data mart.

Importing, Exporting, and Promoting Metadata

About Importing, Exporting, and Promoting Metadata

SAS IT Resource Management metadata objects are managed as all other SAS metadata objects are managed. The SAS Intelligence Platform provides several tools to assist you in copying, promoting, importing, and exporting your metadata. For information about

how metadata can be managed, see the “Using the Promotion Tools” chapter in the *SAS Intelligence Platform: Data Administration Guide*.

Note: SAS IT Resource Management does not support the exporting or importing IT data marts that contain duplicate named items, even if those items are in different folders. If you want to export (and then import) IT data marts that contain duplicate items, you should first export (and then import) the IT data mart without including its dependent objects. After that, as subsequent steps, you can export (and then import) each separate first-level subfolder including its dependent objects.

About Promoting an IT Data Mart

The IT data mart object is unique to SAS IT Resource Management and enables you to perform actions on the whole IT data mart. One of these actions is the ability to promote the IT data mart by means of the import and export tools. These tools enable you to move the definition of an IT data mart from one metadata server to another. For example, you might want to move (or promote) an IT data mart for one of the following reasons:

- to move an IT data mart from a test system to a production system
- to move an IT data mart to a new server when the original host is being replaced by new or different hardware

Note: Both servers must be SAS 9.3 servers.

- to share an IT data mart within or across an organization
- to share user-written or customized adapters across IT data marts or IT Resource Management systems

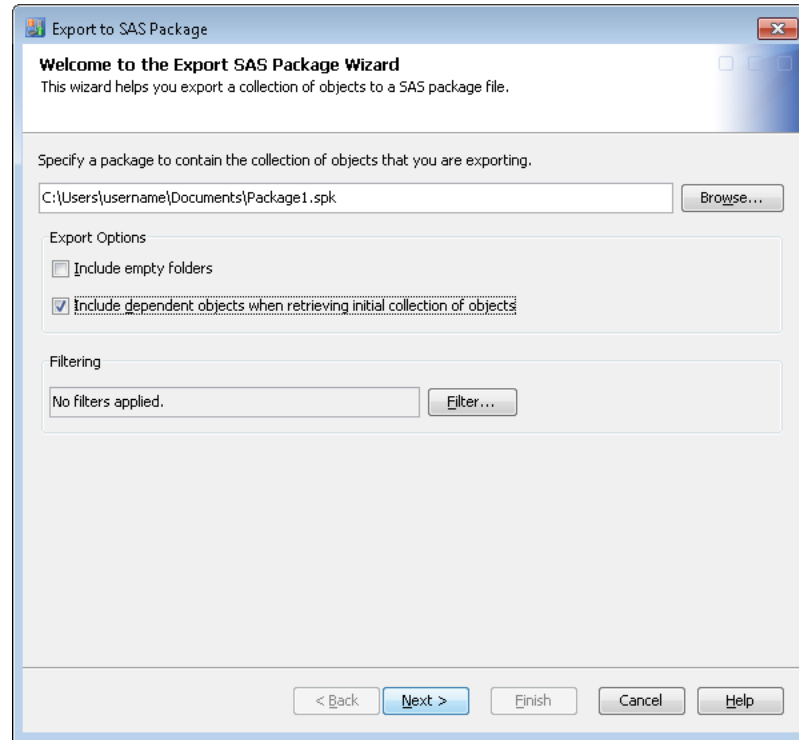
Note: Although you can invoke the import and export wizards from the Folders tree, for best results, invoke these wizards from the IT Data Marts tree.

How to Export an IT Data Mart

To export an IT data mart, perform the following steps:

1. In the IT Data Marts tree, navigate to the IT data mart that you want to move.
2. Right-click the IT data mart and select **Export SAS Package**.

The first page of the Export SAS Package Wizard opens.

Display 4.2 First Page of the Export SAS Package Wizard

3. In the **Specify a package to contain the collection of objects that you are exporting** field, you can accept the default location. The package is called **Package nn .spk**, where nn is a number that is incremented by the value “1” each time the wizard is invoked.

Alternatively, click **Browse** to navigate to another location. Select the location that you want to use and click **OK**.

In addition, if you want to include all the contents of the IT data mart, make sure that **Include dependent objects when retrieving initial collection of objects** is checked.

4. On the **Select Objects to Export** page, all the contents of the selected IT data mart are selected. You can deselect any items that you do not want to export.
5. The **Summary** page displays the objects that will be included in the SAS package. When you are satisfied with your selections, click **Next** to export the IT data mart and its selected objects.

Click **View Log** to see the log file that was created during the export process.

6. Click **Finish** to return to the IT Data Marts tree.

How to Import an IT Data Mart

If the exported IT data mart contains any staging jobs whose rawdata input is a database, the Import wizard might require you to select a library to which the exported library can map. If the library that points to the input database is located within the IT data mart that is being promoted, then you will not be required to specify a new library definition.

In addition, if the exported IT data mart contains any staging jobs whose rawdata input is a database, the Import from SAS Package wizard will require you to select a database server definition. This database server is associated with the library that is specified on

the staging job as rawdata. If you want to work with a database server definition that does not exist, you must create one. For information about how to register a database server, see the “Connecting to Common Data Sources” topic in the *SAS Intelligence Platform: Data Administration Guide*.

TIP If you need to specify either a library or a database server, do so before you invoke the Import from SAS Package wizard.

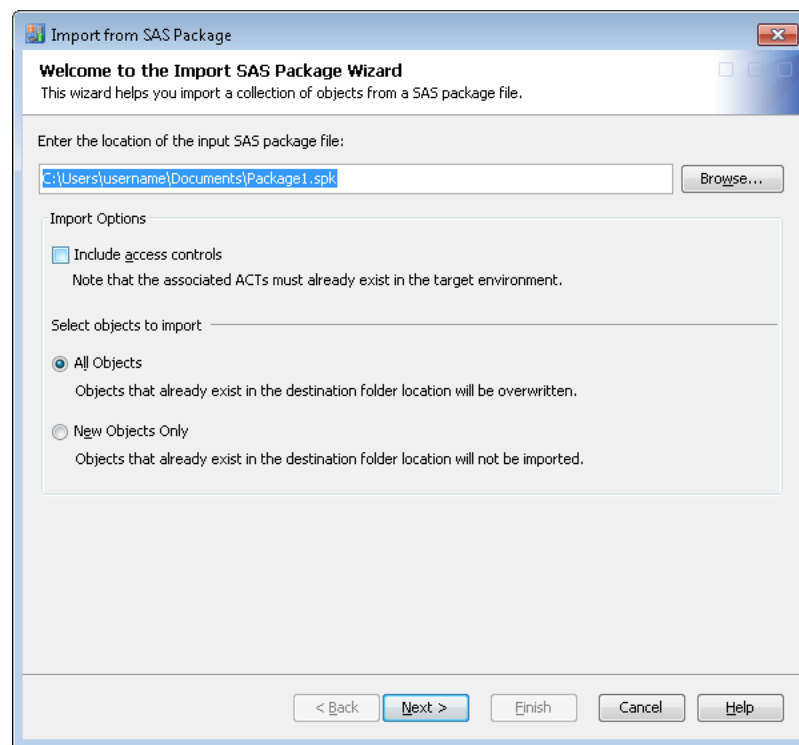
For best results, make sure that the name of the IT data mart that you are importing is unique within the metadata server. If the metadata server already contains an IT data mart with the same name as the one being imported, the name of the IT data mart is automatically changed to “Copy of <original-name>” or, if that name is already in use, “Copy n of <original-name>”.

Note: You might be importing the IT data mart to a machine that is running a different version of SAS IT Resource Management from the one where the source data mart resides. Similarly, you might be importing the IT data mart to a different machine from the one where the source IT data mart resides. In either case, any deployed jobs within the imported data mart on the target system should be deleted and then deployed again. You should not redeploy them.

To import an IT data mart, perform the following steps:

1. Right-click in the white space of the IT Data Marts tree, and select **Import SAS Package**. The first page of the Import from SAS Package Wizard opens.

Display 4.3 First Page of the Import SAS Package Wizard



Note: If you have selected an IT data mart, the **Import SAS Package** selection is not available. The **Import SAS Package** selection is available only if an IT data mart has not been selected.

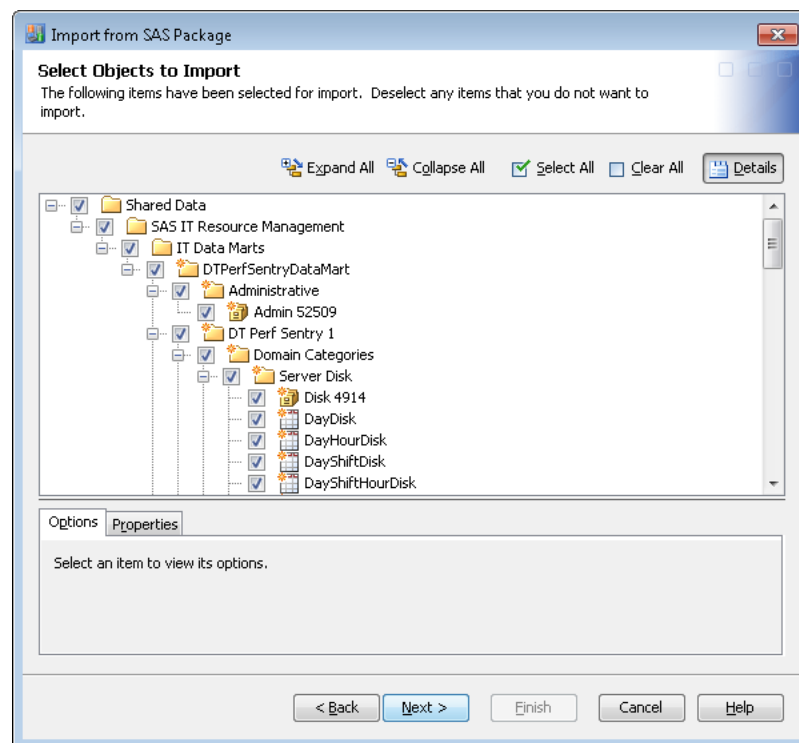
2. On this page, you can accept the default location, which is the location of the most recent exported SAS package.

Alternatively, click **Browse** to navigate to another SAS package. Select the SAS package that you want to use and click **OK**.

You can also choose the following options:

- include the access controls. These access controls govern which users or user groups can access the imported IT data mart.
 - import all objects in the selected SAS package. This option causes objects that exist in the destination IT data mart to be overwritten.
 - import only new objects in the selected SAS package. This option imports only those objects that do not exist in the destination IT data mart.
3. On the second page of the wizard, you can select the objects that you want to import. All items in the SAS package are selected by default. You can deselect those items that you do not want to import.

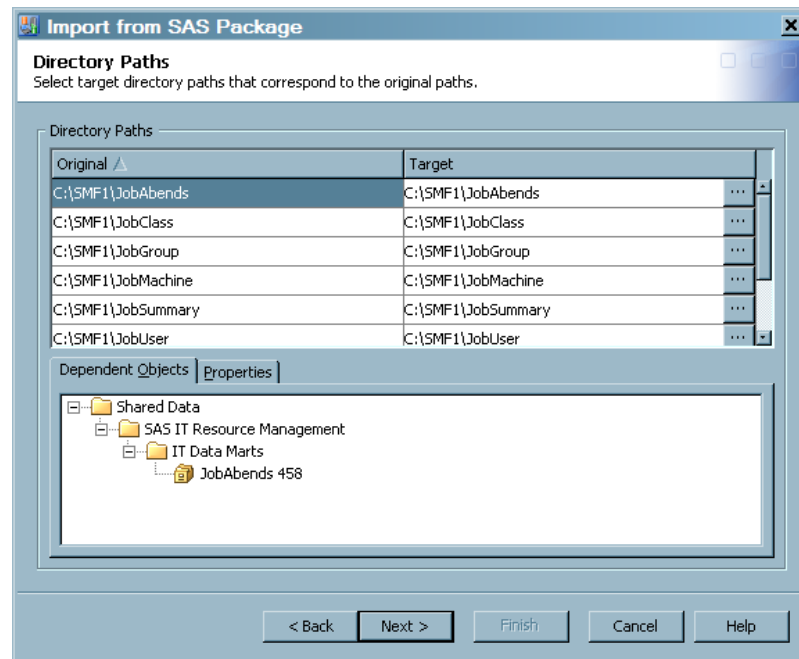
Display 4.4 Select the Objects to Import



4. The next page of the Import from SAS Package Wizard displays a note that explains that you might need to create metadata definitions for some of the objects that you want to import. For example, libraries that are associated with tables must be created before the import process is started. You can create a new library definition with the New Library wizard. To access this wizard from the SAS IT Resource Management client menu bar, select **File** ⇒ **New** ⇒ **Library**. Respond to the prompts of the wizard as needed. In addition, some adapters (such as MS SCOM, SNMP, or VMware) have a raw data source that is a database. If an IT data mart contains that type of adapter, then a database server definition must be created that points to that database. For information about how to create a database server definition, see the “Overview of Common Data Sources” chapter in the *SAS Intelligence Platform: Data Administration Guide* at <http://support.sas.com/documentation/onlinedoc/intellplatform/index.html>. In the topic that discusses your database, you can find the section called “Register the Database Server.”

5. The next page of the wizard prompts you to specify the target application server to which you are importing the IT data mart.
6. The Directory Paths page shows the target locations for the folders of the IT data mart that is being imported.

Display 4.5 Select Target Directory Paths for IT Data Mart Objects



You can accept the default target locations, or you can click the ellipsis (...) to open the Browse dialog box. When you have selected the new target location, click **OK**.

7. The Summary page displays the choices that you specified. Click **Next** to import the IT data mart and the objects that you specified.
8. When the IT data mart has been imported, the Import Complete page is displayed. Click **View Log** to see the log file that was created during the import process.
9. Click **Finish** to return to the IT Data Marts tree.
10. Use your file system management tools to copy the physical files that are associated with the IT data mart from the source to a location on the target server that is to be referenced from the IT data mart. For more information, see *SAS IT Resource Management: Migration Documentation*.

Chapter 5

Adapters

About Adapters	67
What are Adapters?	67
Supported Adapters	68
Adapter Components That Stage Data	75
Working with Template Tables	75
What are Template Tables?	75
Locate Template Tables	76
Differentiate between Template Tables and Staged Tables	77
Template Table Properties	77
Computed Columns in Template Tables	82
Use Template Tables to Create Staged Tables	83
Maintaining Template Tables and Staged Tables Based on Adapter Changes . . .	84
Working with the Maintain Staged Tables Wizard	84
Maintaining Staged Tables in Batch Mode	90
Examples of the Batch Maintain Table Command Line	93
Sample Batch Maintain Table Commands for Windows, UNIX and z/OS Operating Environments	94

About Adapters

What are Adapters?

Adapters consist of specialized software that enables you to process data about IT resources from many diverse data sources. Adapters convert raw data from a given source into a set of organized IT resource information. This organization includes staged tables, aggregated tables, analysis- and report-ready data, and reports that facilitate aggregation and reporting on your IT resources.

Raw performance data is generated by the logging mechanisms that are inherent to IT resources. It can also be created by the Enterprise Systems Management tools that are used to manage the IT infrastructure. SAS IT Resource Management accepts data from many different vendors. The only requirement is that the data must include a timestamp so that it can be categorized and aggregated appropriately for subsequent analysis.

Raw data about an IT resource is represented using measures and formats that are specific to the software that collected it. In order for the data to be processed into meaningful statistics about an enterprise's IT resources, the raw data must be read, interpreted, and possibly converted into a more standard form. The data can then be

loaded into staged tables in preparation for subsequent aggregation and reporting. For each type of IT performance data that is supported, SAS IT Resource Management supplies an adapter that consists of the transformations, template tables, and code that are necessary to stage, aggregate, and report on that IT performance data.

For most adapters, by means of the Adapter Setup wizard, SAS IT Resource Management also supplies Aggregation, Information Map, and Gallery transformations. These transformations create the aggregation tables, information maps, and reports that enable you to convert raw data from various sources into meaningful intelligence about your IT resources. If these transformations are not created by the Adapter Setup wizard, then they can be created manually. To do so, see “Working with Aggregation Tables” on page 140 and “Overview of Specifying an Information Map Transformation” on page 212. For more information about using Gallery transformations and reporting jobs, see the *SAS IT Resource Management: Reporting Guide*.

Adapters stage the data by means of transformations that are executed by the SAS Data Integration Studio jobs. For information about how to work with SAS Data Integration Studio, see the documentation that is available for this product at <http://support.sas.com/documentation/onlinedoc/etls/index.html>.

Supported Adapters

The name of an adapter represents the data source from which SAS IT Resource Management gathers raw data.

Note: All SAS IT Resource Management adapters are supported on all SAS IT Resource Management operating environments unless there is a technical limitation that is associated with the raw data (log file) sources. For example, if the raw data resides on a device that is not compatible with the architecture of the host system on which the SAS IT Resource Management server is installed, then that raw data file might not be usable as input unless further processing is performed.

Table 5.1 Table of Supported Adapters

Adapter Name	Description	Latest Verified Version of the Data Source
ASG TMON2CIC	<p>ASG TMON for CICS TS for z/OS (formerly from Landmark) monitors the CICS Transaction Server (TS) and provides resource consumption measurements by tracing each transaction's performance by CICS event and related unit of work.</p> <p>Note: CICS data can be produced by either ASG TMONCICS or as a subtype of SMF.</p>	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
ASG TMONDB2	ASG TMON for DB2 (formerly from Landmark) works with the SQL Analyzer to manage DB2 applications and critical DB2 resources by providing a single view of DB2 data-sharing group performance. ASG TMON for DB2 tracks DB2 buffer pools, lock contention and serialized resource usage, transaction statistics from ASG TMON for CICS to ASG TMON for DB2, and dynamic and static SQL calls. ASG TMON for DB2 provides a view for all performance statistics by plan, package, and statement, including CPU time, I/O analysis, lock activity, and buffer pool utilization by SQL statement.	MXG 29.04*
BMC Mainview IMS	BMC Mainview for IMS (formerly IMF) supplies transaction information such as CPU usage, number of I/Os, and response time. When you collect BMC Mainview for IMS data instead of IBM IMS data, BMC recommends that you collect a duration of one hour of consistently high activity, starting at the beginning and ending at the end of an RMF interval.	MXG 29.04*
BMC Perf Mgr	BMC Performance Manager for Servers (formerly BMC PATROL) collects UNIX and Windows enterprise system management data.	v3.7.30
CA TMS	The Computer Associates CA 1 Tape Management product (formerly TMS) controls and protects tape data sets and volumes in z/OS environments.	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
DT Perf Sentry	<p>Demand Technology Performance Sentry (formerly NTSMF) collects data from Windows systems and for Windows server applications such as Exchange, SQL Server, and IIS.</p> <p>SAS IT Resource Management also supports a separate Demand Technology Performance Sentry adapter that is z/OS based. The Demand Technology Performance Sentry with MXG adapter that is z/OS based is for raw data sources that are supported through MXG.</p>	v3.1.4
DT Perf Sentry with MXG	<p>Demand Technology Performance Sentry (formerly NTSMF) with MXG collects data from Windows systems and for Windows server applications such as Exchange, SQL Server, and IIS. (SAS IT Resource Management also supports a separate Demand Tech Perf Sentry adapter that is not MXG based.)</p>	MXG 29.04*
HP Perf Agent	<p>HP Performance Agent (formerly OpenView Performance Agent) captures enterprise systems management measurement data for IT resources from Windows systems and many UNIX variants such as HP-UX, Sun OS/Solaris, IBM AIX, Tru 64 UNIX, and Linux.</p>	v5.0
HP Reporter	<p>HP Reporter software (formerly OpenView Reporter) captures HP Performance Agent measurement data that is stored in a relational database via SAS/ACCESS to Oracle or ODBC.</p>	v3.70
IBM AS400	<p>IBM OS/400 Collection Services</p> <p><i>Note:</i> The staging code that is required for the IBM AS400 data source is not automated by a SAS IT Resource Management adapter. Instead, SAS IT Resource Management supplies a User-written Staging transformation and templates that contain table and column metadata for this adapter for use with user-written staging code.</p>	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
IBM DCOLLECT	<p>IBM DFSMS Data Collection Facility obtains the following Mainframe Storage System data:</p> <ul style="list-style-type: none"> • active data sets • VSAM data set information • volumes • inactive data • migrated data sets • backed-up data sets • capacity-planning data such as the following: <ul style="list-style-type: none"> - DASD capacity planning: IBM DCOLLECT provides information and statistics for volumes that are managed by DFSMSHsm (ML0 and ML1). - tape capacity planning: IBM DCOLLECT provides statistics for tapes that are managed by DFSMSHsm. • SMS configuration information 	MXG 29.04*
IBM EREP	<p>IBM Environmental Record Editing and Printing Program processes the error records from the MVS, VM, and VSE operating systems to produce formatted reports. These IBM EREP reports can show the status of the entire installation, an I/O subsystem, or an individual device, depending on which report you request. IBM EREP reports can vary in format, depending on the following report types:</p> <ul style="list-style-type: none"> • System summary reports show error data in summary form. • Trend reports show error data by daily totals. • Event history reports show error data in a time sequence by occurrence. 	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
IBM IMS	<p>IBM Information Management System (IMS) is the transaction and hierarchical database management system from IBM.</p> <p><i>Note:</i> The staging code that is required for the IBM IMS data source is not automated by a SAS IT Resource Management adapter. Instead, SAS IT Resource Management supplies a User-written Staging transformation and templates that contain table and column metadata for this adapter for use with user-written staging code.</p>	MXG 29.04*
IBM SMF	<p>IBM System Management Facility collects and records system and job-related information that is used for the following purposes:</p> <ul style="list-style-type: none"> • billing users • reporting reliability • analyzing the configuration • scheduling jobs • summarizing direct access to volume activity • evaluating data set activity • profiling system resource use • maintaining system security <p><i>Note:</i> RMF is a subtype of SMF records. IBM RMF is the IBM strategic product for z/OS performance measurement and management. IBM RMF is the base product that collects performance data for z/OS and Sysplex environments to monitor system performance. Because IBM RMF data is a subtype of IBM SMF data, the IBM SMF adapter can be used with IBM RMF data without requiring additional customizations.</p>	MXG 29.04*
IBM TPF	<p>The IBM Transaction Processing Facility operating system works with application programs to process transactions in a real-time environment. The IBM TPF system is designed for businesses and organizations that have large networks and high volumes of online transactions.</p>	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
IBM VMMON	IBM VM Monitor Collection product collects performance information that is associated with VM, including measurements for user activity, processor storage, I/O, and applications.	MXG 29.04*
MS SCOM	Microsoft System Center Operations Manager monitors Windows systems, Microsoft software, and applications to provide a view of the health of an organization's Windows environment.	v2007
SAP ERP	<p>Enterprise Resource Planning (ERP) Software from SAP (formerly SAP R/3), via the SAS IT Management Adapter for SAP, collects SAP ERP and SAP Business Warehouse (BW) workload performance measurements.</p> <ul style="list-style-type: none"> For SAP ERP, measurements are read from the SAP Statistic File (stat file). For SAP BW, the measurements that communicate the amount of computer resources that were consumed to populate and query BW cubes are read. These measurements are then used to supplement BW server performance statistics. 	v7.0
SAR	System Activity Reporter is a logging mechanism that is native to most UNIX and Linux variants. SAR captures the contents of cumulative system activity counters.	<p>Dependent on the operating system.</p> <p>For information about specific operation system release versions, see “Accessing Raw Data for SAR” in Appendix 2, “Data Sources Supported by SAS IT Resource Management Adapters” in SAS IT Resource Management: Administrator's Guide.</p>
SNMP	<p>Simple Network Management Protocol is a component of the Internet Protocol Suite.</p> <p>The SNMP adapter also supports gathering SNMP data from round-robin databases (RRDs) that are used in many network management tools.</p>	Not applicable

Adapter Name	Description	Latest Verified Version of the Data Source
VMware vCenter	VMware vCenter Server is an operational console for VMware and includes the VMware vCenter database that serves as the raw data source for this SAS IT Resource Management adapter. The vCenter database contains VMware performance and utilization measurements for both the virtual machines and the physical servers on which they run, including detailed CPU, memory, storage, and I/O metrics. The vCenter relational database is accessed via SAS/ACCESS to Oracle, DB2, or ODBC.	v4.0
VMware Data Acquisition	<p>VMware Data Acquisition is a transformation that can extract topology and configuration information from the VMware vCenter database more frequently than once a day. The vCenter relational database is accessed via SAS/ACCESS to Oracle, DB2, or ODBC.</p> <p><i>Note:</i> This is not an adapter. However, a staging transformation is available for this source.</p>	v4.0
Web Log	The Web Server Log adapter reads and analyzes Web logs that conform to the Common and W3C Extended Log Formats (CLF and ELF.)	Not applicable

* SAS IT Resource Management uses MXG for this adapter.

Note: The specific software versions of the raw data sources in this table have been used to verify each SAS IT Resource Management 3.21 adapter. This information is not an exclusive list of all data source versions that are supported. Earlier versions of some data source software might be supported. Future releases of some data source software might be compatible and supported by adapters for SAS IT Resource Management 3.21. For information about supporting a data source version that is not included in the previous table, contact your SAS Installation Representative or on-site SAS support personnel.

For data sources (and for raw data sources) that are user-written, staging code is supported by the wizards of SAS Data Integration Studio, and the transformations that are supplied by SAS IT Resource Management. The New Table wizard of SAS Data Integration Studio can create the transformations that can stage raw data from any input source.

Adapter Components That Stage Data

SAS IT Resource Management supports adapters that are essential for extracting, standardizing, and staging data. These adapters include the following three components to stage data:

template tables

These adapter-specific models are used when transforming and loading raw data into SAS data sets called staged tables. The values in template tables are supplied by SAS IT Resource Management for specific IT data sources and thus cannot be changed. However, you can view a template table's properties such as associated columns and related metadata. For additional information about template tables, see [“Working with Template Tables” on page 75](#).

staging transformations

These transformations generate the SAS code that extracts raw data, transforms it to a standardized format, and then loads it into staged tables. Staging transformations include specific staging parameters such as the format of raw data, the location of the data, and how the data will be handled during the staging process. For example, a staging transformation might specify the network path to the raw data, how that data is delimited in its source format, and how duplicate data might be handled when it is staged.

SAS IT Resource Management provides several ways that you can customize and configure the way that your data is staged by a staging transformation. For example, you can use the Adapter Setup wizard to create and configure staging transformations for a supplied adapter. You can create staging transformations manually by using the User-written Staging transformation. You can also modify the staging parameters that are specified in a staging transformation after it is created. For more information about staging transformations, see [“Working with Staging Transformations” on page 107](#).

staged tables

These SAS data sets contain the data that has been extracted from an input data store and rendered into a format that is suitable for further transformation and generation into reports. The structure and metadata for staged tables are based on corresponding template tables that SAS IT Resource Management provides for specific adapters and domain categories. For more information about staged tables, see [“Working with Staged Tables” on page 112](#).

Working with Template Tables

What are Template Tables?

Template tables are representations of adapter-specific tables that are supplied by SAS IT Resource Management for specific IT data sources. Template tables are models that are used to create staged tables. These staged tables are, in turn, populated by staging transformations when loading and transforming raw data.

Each template table provides the metadata for all regular columns, computed columns, and other attributes (formats, formulas, source code for computed columns, and so on) that a staged table can include when it is created from a template table. The resulting

In the previous image from the Folders tree view, the template table PCSDSK is in the **Disk** subfolder and is therefore related to the disk performance data from the HP Reporter adapter. Using this information, you can determine that this template table provides metadata for staging the data that relates to PCS disk metrics such as disk performance and disk space usage. Other template tables that relate to other facets of disk performance are also in the **Disk** subfolder.



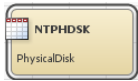
Note: Domain categories also surface in the Staged Table wizard and the Adapter Setup wizard for creating staged tables.

Differentiate between Template Tables and Staged Tables

Staged tables have the same name as the template tables on which they are modeled. For this reason, identifying an object as a template table or a staged table when working within SAS Data Integration Studio can be confusing.

The following table shows the differentiating factors that enable you to distinguish a template table from a staged table.

Table 5.2 Distinguishing Template Tables from Staged Tables

Differentiator	Template Tables	Staged Tables
Icon	Template tables are indicated by the  icon in the tree views and the Staged Table wizard.	Staged tables are indicated by the  icon in the tree views.
Location	In the Folders tree view, template tables are in the Products ⇒ SAS IT Resource Management ⇒ Adapters ⇒ <adapter name> ⇒ IT Template Tables ⇒ <domain category> subfolders.	In the Inventory tree view, staged tables are in the Table folder. In the IT Data Marts tree view, staged tables are in the same folder as the corresponding staging job (unless you designated otherwise when creating the staged tables). Staged tables are also stored in the same IT data mart subfolders as their respective staging jobs.
Job	A template table never appears in a folder of the IT Data Marts tree view or in the process flow diagram of a job.	Each staged table appears as an object in the process flow diagram of a staging job. The following image is an example of how a staged table looks in a process flow diagram. 

Template Table Properties

About Template Table Properties

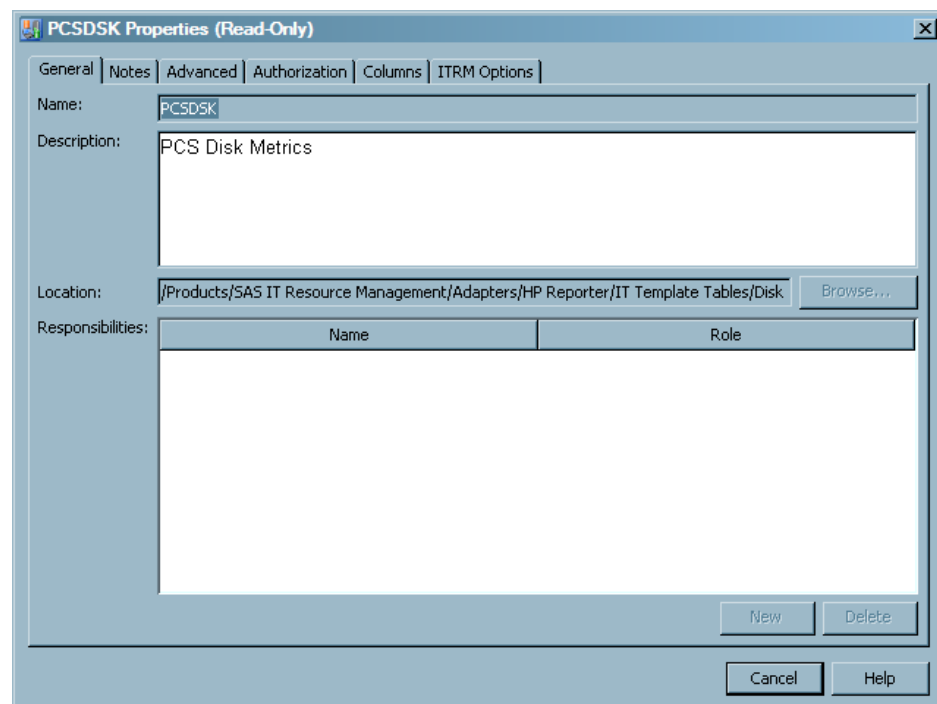
Template tables are distinct tables that are defined to the metadata for SAS IT Resource Management. You cannot modify the properties of template tables because these

predefined tables include values and attributes that make up the intelligence for processing raw IT data based on industry research. You can, however, view the properties of template tables to determine the template tables from which you want to create staged tables for a given adapter.

To view the properties of a template table, perform the following steps:

1. In the Folders tree view of SAS IT Resource Management, expand **Products** ⇒ **SAS IT Resource Management** ⇒ **Adapters**. Then expand an adapter subfolder, the **IT Template Tables** subfolder, and one or more of the domain category subfolders to see the corresponding template tables.
2. Right-click a template table and select **Properties**. A Properties dialog box appears and enables you to view the template table's metadata, such as table name, column metadata, and external properties.

Display 5.2 Properties Dialog Box for the PCSDSK Template Table



The Properties dialog box contains the following tabs:

- The **General** tab displays information that identifies the template table, such as name, description, folder location, and responsibilities.
- The **Notes** tab displays notes or documents that are associated with the template table.
- The **Advanced** tab displays additional metadata attributes about the template table, such as metadata ID and filename.
- The **Authorization** tab displays users and groups that have various permissions related to template table.
- The **Columns** tab displays the metadata for each column in the template table. This tab includes information such as the name and description for each column, the expressions used for any computed columns in the template table, and the type classification for each column in the template table. For more information about the **Columns** tab for template tables, see [“Columns in Template Tables” on page 79](#).

- The **ITRM Options** tab identifies the adapter that is associated with the template table as well as any predecessor tables, external names, and data model versions that correspond with the template table. For more information about **ITRM Options** tab for template tables, see “[ITRM Options for Template Tables](#)” on page 81.

Note: You can click **Help** in the Properties dialog box for more information about each tab.

Columns in Template Tables

The **Columns** tab of a template table's Properties dialog box enables you to view the specific column metadata for the table. You can use this information to view the types of data that the template table includes, understand how computed columns are formulated, and determine whether you want to use the template table for creating a staged table.

Display 5.3 Sample Columns Tab from the Properties Dialog Box of the PCSDSK Template Table

PCSDSK Properties (Read-Only)

General | Notes | Advanced | Authorization | Columns | ITRM Options

#	Name	Expression	Description	Type	Length	Informat	Format	External Name	Recommended
114	DSKVRCP		DiskSpaceCapacityVirtual...	Numeric	6 (None)	NLNUM16.2	GBL_DISK_VIRTUAL_CA...	No	
115	DSKVRRT		DiskVirtualMemoryRead...	Numeric	6 (None)	NLNUM16.2	BYDSK_VM_READ_RATE	No	
116	DSKVRT		DiskVirtualMemoryWrite...	Numeric	6 (None)	NLNUM16.2	BYDSK_VM_WRITE_RATE	No	
117	DSKWNT		DiskTransferWriteKB	Numeric	6 (None)	NLNUM16.2	BYDSK_PHYS_WRITE_BYTE	Yes	
118	DSKWRAT		DiskWriteTransferRateKB	Numeric	6 (None)	NLNUM16.2	BYDSK_PHYS_WRITE_BY...	Yes	
119	DSKYEAR		RecordYear	Numeric	6 (None)	NLNUM16.0	YEAR	No	
120	DURATION		Duration	Numeric	6 (None)	TIMES.		Yes	
121	GLBDSKI		DiskIndex: Number	Numeric	6 (None)	(None)	GBL_DISK_INDEX	No	
122	LSTPDAT		LastProcessDate	Numeric	8 (None)	NLDATM18.	ProcessDate	Yes	
123	MACHINE		Machine	Character	32 (None)	\$32.		Yes	
124	OS		OpSys	Character	16 (None)	(None)		Yes	
125	DATETIMEFIFTE...	intnx('DTMINUTE15', dat...	DatetimeFifteenMinute	Numeric	8 (None)	NLDATM18.		Yes	
126	DATETIMEFIVEMIN	intnx('DTMINUTES', datet...	DatetimeFiveMinute	Numeric	8 (None)	NLDATM18.		Yes	
127	DATETIMEONEMIN	intnx('DTMINUTE1', datet...	DatetimeOneMinute	Numeric	8 (None)	NLDATM18.		Yes	
128	DATETIMETENMIN	intnx('DTMINUTE10', dat...	DatetimeTenMinute	Numeric	8 (None)	NLDATM18.		Yes	
129	DATETIMETHIRT...	intnx('DTMINUTE30', dat...	DatetimeThirtyMinute	Numeric	8 (None)	NLDATM18.		Yes	
130	DATETIMETWEN...	intnx('DTMINUTE20', dat...	DatetimeTwentyMinute	Numeric	8 (None)	NLDATM18.		Yes	
131	DAYDATE	intnx('DAY', datepart(dat...	DayDate	Numeric	8 (None)	NLDATM10.		Yes	
132	DAYOFMONTH	day(datepart(datetime))	DayOfMonth	Numeric	8 (None)	2.		Yes	
133	DAYOFWEEK	weekday(datepart(datet...	DayOfWeek	Numeric	8 (None)	WKD.		Yes	
134	DSKFRWATE	sum(DSKFRD, DSKFWT)	DiskFilesystemIoRate	Numeric	8 (None)	NLNUM16.2		Yes	
135	DSKIDLEPCT	sum(100, -(DSKUTL * 100))	DiskIdlePct	Numeric	8 (None)	NLNUM16.2		Yes	
136	DSKLWRATE	sum(DSKLRD, DSKLWT)	DiskLogicalIoRate	Numeric	8 (None)	NLNUM16.2		Yes	
137	DSKTRFMB	sum(DSKTRFMB)	DiskTransferRateMB	Numeric	8 (None)	NLNUM16.2		Yes	

Cancel | Help

The following list describes the information that you can view for each column in a template table. Each row on the **Columns** tab shows the metadata for a single column of the template table. You cannot modify this information because the predefined metadata includes values and attributes that make up the intelligence for processing raw IT data based on industry research.

#

specifies the column's ordinal position in the grid on the **Columns** tab. Note that the ordinal position is not necessarily the same as the alphabetic sequence of variable names.

Name

specifies the name of the column in the template table.

Expression

specifies the combination of functions and mathematical operations that are used to derive a value for the column. This field is blank unless the column is a computed column. For more information about computed columns in template tables, see “[Computed Columns in Template Tables](#)” on page 82.

Description



provides a description of the data that populates the column. The description provides more information about the type of data that the column includes.

The column descriptions in template tables are propagated to staged tables and used as labels in information maps and aggregations. Therefore, column descriptions for template tables are standardized to be concise, easily understood, and efficiently integrated into other objects and SAS products such as SAS Enterprise Guide and SAS Information Map Studio. Standardizing column descriptions and label names ensures that equivalent data is labeled the same across all adapters. Standardization also reduces the need for editing label content in reporting clients.

Here are some of the ways that column descriptions in template tables have been standardized to be effective labels:

- Character length is reduced by removing spaces between words and capitalizing the first letter of each word or term within the description (such as *ManagementClassLength*).
- Character length is typically reduced by removing special characters, punctuation, and redundant words that do not enhance the meaning of the data.
- Words that indicate performance area or domain category (such as CPU, processor, disk, memory, and network) are typically at the beginning of the description.
- Words that indicate data interpretation and normalization (such as rate, bytes, and seconds) are at the end of the description.
- Words are abbreviated if a clear alternative is available (such as *Average* to *Avg*).

Type

specifies the data type of the column. Valid types are character and numeric. You can also identify column types by looking at the Name variable. A round  icon to the left of the name indicates that the data in the column is numeric. A pyramid  icon to the left of the name indicates that the column contains character data.

Length

specifies the maximum length of the data in the column.

Informat

specifies the SAS informat that is used (if needed) to read values for the selected column. Informats are required for reading fixed-width data.

Format

specifies the SAS format that is used to write values for the selected column. Template table columns are in NLS format when applicable. NLS formats are indicated by **NL** as the first two characters in the format column. For more information about NLS formatting, see [“National Language Support \(NLS\)” on page 12](#).

External Name

specifies the name by which the staging code maps columns in a staged table to their corresponding variables in the raw data. External names are used by both staged tables and template tables. However, not all columns in a staged or template table have a predetermined external name.

Note: For MXG adapters, the external name is usually the name of the MXG variable. For the DT Perf Sentry adapter, the external name is the name of the DT Perf Sentry metric.

Recommended

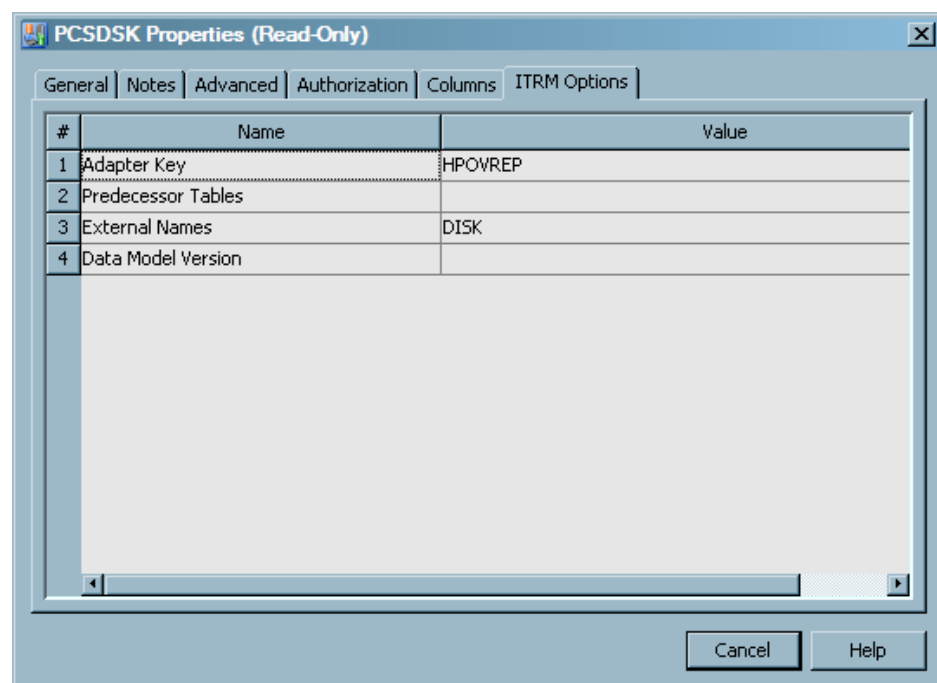
specifies whether the column provides data that SAS IT Resource Management has predefined as critical for reporting and analysis in the performance area of the template table. The Staged Table wizard enables you to use only the recommended columns from a template table when creating a corresponding staged table. This enables you to create staged tables that include only the columns that are recommended and thus not spend resources staging and analyzing data that is likely not relevant. For more information about using the Staged Table wizard, see [“Create Staged Tables” on page 114](#).

The Adapter Setup wizard also uses the recommended columns of a template table to create various levels of output. For example, when using the Adapter Setup wizard, you can choose staging only, minimal, typical, or full output. Then the wizard creates the appropriate level of output based on your selection and the columns recommended for the corresponding tables. For more information about the Adapter Setup wizard, see [“About the Adapter Setup Wizard” on page 177](#).

ITRM Options for Template Tables

The **ITRM Options** tab of the Properties dialog box for a template table includes information such as what tables are staged before the template table and how the original data collector identifies the data that is used for the template table.

Display 5.4 Sample ITRM Options Tab from the Properties of the PCSDSK Template Table



The **ITRM Options** tab includes the following parameters:

Adapter Key

specifies the adapter that is associated with the template table. For adapter names, SAS IT Resource Management uses the latest names that are approved by respective third parties. However, due to the constraints of character limitations in software code and organizations changing the names of their products and software, SAS IT Resource Management uses an abbreviation called an adapter key to represent an adapter in the software if necessary. The adapter key is for reference by SAS

Technical Support and does not require any action from the SAS IT Resource Management user.

Note: In some cases, the adapter key does not match the current name of the adapter but instead reflects the former name of the adapter. For example, the adapter key for the BMC Perf Mgr adapter is PATROL, the former name of this adapter.

Predecessor Tables

specifies other tables that are required to be staged in order to produce this table. For example, one of the staged tables (named SAR) for the SAR adapter is a compilation of many SAR tables (predecessor tables). When the SAR table is staged, SAS IT Resource Management programmatically stages the various predecessor tables that make up the resulting SAR staged table so that you do not have to stage each predecessor table individually. If a template table does not require staging predecessor tables, this parameter is blank.

External Names

specifies the names that the original data collector (or other software) uses to identify the data that serves this template table. External names are the means by which the staging code maps a SAS IT Resource Management template table to a table in the raw data. For example, a template table that is named “NTPHDSK” in SAS IT Resource Management might have an external name such as “Physical Disk” in the raw data file. The staging code uses this information to determine that the Physical Disk data in the raw data file belongs in the staged table that is modeled from the NTPHDSK template table.

Note: The staging code for some adapters does not use external names. For example, the table name in the raw data might be the same as the table name in SAS IT Resource Management.

The external names for template tables are similar in concept to the external names for staged tables. For more information about using external names for staged tables and columns in staged tables, see [“Understanding External Names of Staged Tables and Columns” on page 127](#).

Data Model Version

specifies the version of the template table. This is the version from which the staged table is created or updated through the Maintain Staged Tables wizard. For more information about the Maintain Staged Tables wizard, see [“About the Maintain Staged Tables Wizard” on page 84](#).

Computed Columns in Template Tables

Computed columns in a table contain data values that are derived from other columns in the table. Template tables contain computed columns for data that is not directly provided by the raw data source but the data might be useful for reporting or promoting consistency across various data sources.

The following table demonstrates how template tables might use computed columns for a data source for consistency across various data sources. In this example, Data Source 1 includes three columns: Column A, Column B, and Column C, which is the total of Column A and B. Data Source 2 provides only two columns: Column A and Column B. For Data Source 2, a template table might compute the total of Column A and Column B so that the data for Data Source 2 is as comprehensive as Data Source 1.

Table 5.3 Example of a Computed Column in a Template Table

Data Source 1	Data Source 2	Template Table for Data Source 2
Column A	Column A	Column A
Column B	Column B	Column B
Column C (equal to Column A + Column B)		Column C (a computed column equal to Column A + Column B)

Note: A computed column should not be used as a source of calculations for other computed columns. Using a computed column as a source for a computation can be problematic due to the order of computations that are necessary to derive the final value.

To find a computed column in a template table and view the formula that it uses, click the **Columns** tab of a template table's Properties dialog box. All computed columns contain an expression in the **Expression** column of the grid.

Here are some of the reasons for which SAS IT Resource Management might use computed columns in a template table:

- Create datetime derivations such as changing week date and shift values to your local environment.
- Provide normalizations such as converting bytes to kilobytes, seconds to milliseconds, and percentages between 0 and 1 to percentages between 0 and 100.
- Total paired columns to provide a summed value that includes input and output values, or received and sent values.
- Derive opposite values by using a percentage-based computation that typically subtracts one value from another. For example, a computed column might subtract the number of used components from the number of allocated components to determine the number of components that are still available.
- Standardize performance event count to total events across various classifications such as batch jobs, TSO sessions, and started tasks.
- Create class columns when a table contains an instance column that can contain either a global or an individual instance of a performance metric. In this case, a computed column such as CpuId or DiskId is created and set to an individual instance value, or it is set to a blank value for global instances. These new class columns can then be used as targets of filters to divide global and individual instances into appropriate aggregations.
- Convert raw counts to rates.
- Calculate response time counts and percentages.

Use Template Tables to Create Staged Tables

The Staged Table wizard and Adapter Setup wizard enable you to create staging jobs that use template tables to create corresponding staged tables. When executed, the staging jobs load the source data into the staged tables. The staged tables can be used as

input to an Aggregation transformation or other SAS Data Integration Studio transformations that are used for analysis and reporting.

SAS IT Resource Management enables you to select the template tables that you want to use or accept the template tables that are selected programmatically for a given adapter.

- To select the template tables that you want to use, create a staging job manually. For more information about how to create a staging job, see [“Add a Staging Transformation to a Staging Job” on page 109](#). For more information about how to create a staged table, see [“Create Staged Tables” on page 114](#).
- To have template tables preselected programmatically for a staging job, run the Adapter Setup wizard for a supported adapter. The template tables that the Adapter Setup wizard uses are predetermined by SAS IT Resource Management to include the data that is most relevant and appropriate for the level of analysis and reporting that you select. For more information about the template tables that are preselected by the Adapter Setup wizard for each adapter and domain category, see [“Overview of the SAS IT Resource Management Data Model ” on page 459](#).

Maintaining Template Tables and Staged Tables Based on Adapter Changes

Working with the Maintain Staged Tables Wizard

About the Maintain Staged Tables Wizard

The data collectors on which template tables are based might change over time. These changes can require new columns in a template table or revisions to existing columns, such as label or character length. Therefore, when SAS IT Resource Management provides a new software release or update, the definitions for some template tables might also change. After a software update is installed, a preexisting staged table can differ from the new version of its corresponding template table.

The Maintain Staged Tables wizard enables you to perform the following tasks with software updates to template tables:

- View updates that can be propagated to existing staged tables.
- Commit updates to the corresponding staged tables within an IT data mart.
- Commit updates to the corresponding staged tables across all IT data marts.
- Add new columns and update existing columns to the corresponding staged tables.
- Commit only updates that add new columns to the affected staged tables.
- Commit only updates to the recommended columns in the corresponding staged tables.

If there are no staged tables to maintain, the HTML report is not produced.

Note: After operations with the Maintain Staged Tables wizard, you must redeploy any staging jobs that were updated. (Aggregation jobs do not require redeployment because the wizard does not affect the metadata of the aggregated tables.)

SAS IT Resource Management identifies affected template tables and their associated staged tables by the **DataModelVersion** field for each table. For example, a template table and a staged table are eligible for impact analysis and updates if the current

DataModelVersion field of the table does not match the version that is associated with the software update.

For more information about using the Maintain Staged Tables wizard to maintain staged tables within an IT data mart, see [“Using the Maintain Staged Tables Wizard to Update Staged Tables within an IT Data Mart” on page 85](#). For more information about using the Maintain Staged Tables wizard to maintain staged tables across all IT data marts, see [“Using the Maintain Staged Tables Wizard to Update Staged Tables across All IT Data Marts” on page 90](#). For more information about maintaining staged tables using a batch process, see [“Maintaining Staged Tables in Batch Mode” on page 90](#).

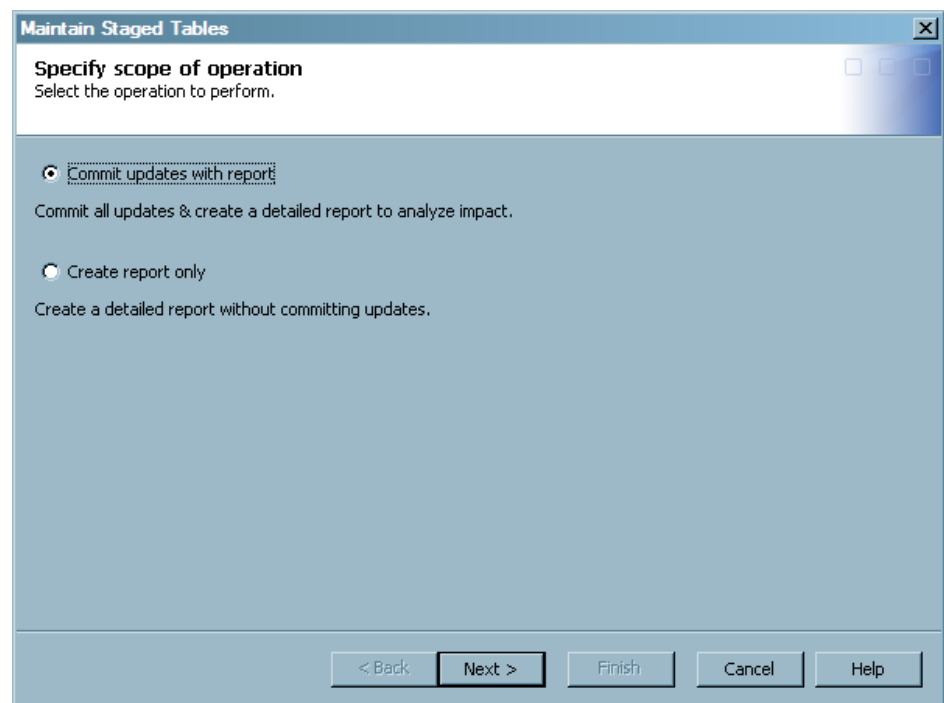
Using the Maintain Staged Tables Wizard to Update Staged Tables within an IT Data Mart

After you install software updates for SAS IT Resource Management, you can use the Maintain Staged Tables wizard to see whether any software changes to the template tables affect the staged tables that you use and propagate these template table changes to their corresponding staged tables. To do so, perform the following steps:

1. In the IT Data Marts tree view of SAS IT Resource Management, locate the IT data mart that contains the staged tables for which you want to view or commit updates.
2. Right-click the IT data mart and select **Maintain Staged Tables** to open the wizard.
3. Specify the scope of the operation that you want to perform with the Maintain Staged Tables wizard.

Note: If there are no staged tables to maintain, the HTML report is not produced.

Display 5.5 Specify Scope of Operation Page in the Maintain Staged Tables Wizard



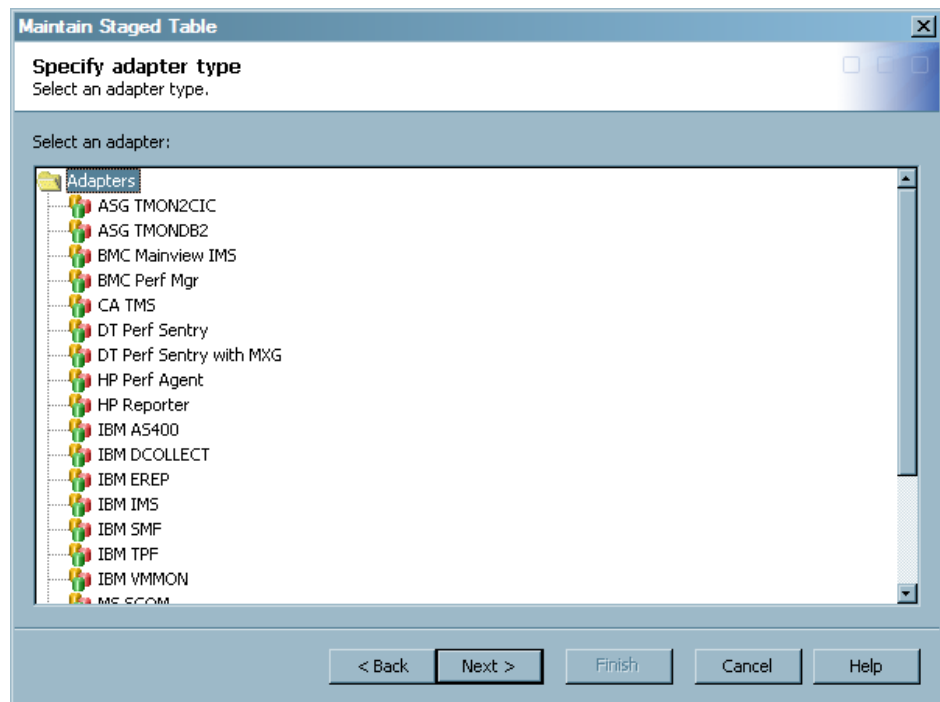
- a. Select **Commit updates with report** to view a report of the template table changes that affect your staged tables and then commit the changes to the corresponding staged tables.
- b. Select **Create report only** to view only a report of the template table changes that affect your staged tables but not commit the changes in this instance of the wizard.

c. Click **Next**.

4. Select an adapter for which you want to analyze the impact of table revisions that are available.

Note: This page is not displayed when accessing the wizard from a specific template table to commit updates across all IT data marts. For more information about using the Maintain Staged Tables wizard to maintain staged tables across all IT data marts, see [“Using the Maintain Staged Tables Wizard to Update Staged Tables across All IT Data Marts”](#) on page 90.

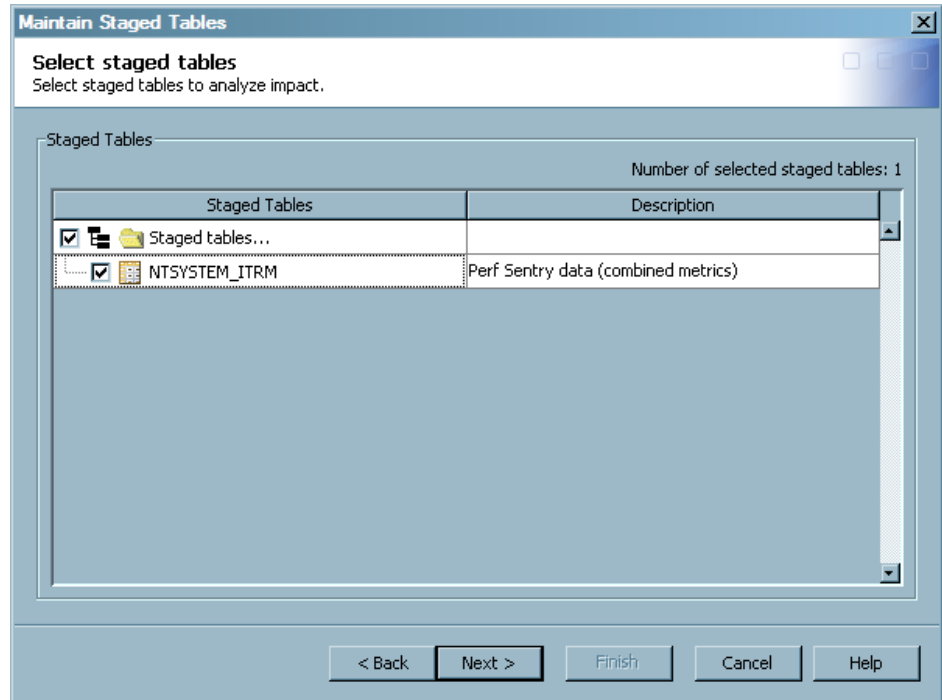
Display 5.6 Specify Adapter Page in the Maintain Staged Tables Wizard



- a. Click the adapter for which you want to analyze updates. The list of adapters includes all of the adapters that SAS IT Resource Management supports. However, you must select an adapter that is used in the IT data mart in order to complete the wizard successfully.

Note: You can select only one adapter on this page and must perform updates for another adapter by using a separate instance of the wizard.

- b. Click **Next**. If a message states that there are no staged tables that need maintenance, then there are no staged table updates required based on your selections and the latest software updates. You can close the message box and stop here. Otherwise, continue to the next step.
5. Select the staged tables for which you want to view changes or commit updates.

Display 5.7 Select Staged Tables Page in the Maintain Staged Tables Wizard

- a. Select the staged tables for which you want to analyze updates. The list of staged tables includes only the staged tables that are associated with the selected adapter.

A tooltip is available for each staged table in the list. The tooltip shows the metadata path of the table in order to help distinguish between tables, especially tables that might have duplicate names.

Note: If there are no updates to the template tables for the selected adapter, then the wizard displays a message to that effect.

- b. Click **Next**.
6. Compare table updates with the previous value of the tables and columns and specify the commit options.

Display 5.8 Impact Analysis Page in the Maintain Staged Tables Wizard

- a. Expand the folders in the grid to view the version of the tables and the table columns that have been modified or added to the corresponding template tables. The rows in the grid include the table name, version, description, or columns that are affected by the update. This comparison list also shows previous values before the update and the corresponding new values after the update.

New columns that are added to the template table display the previous value as blank because the column did not exist in the previous version. New columns are also preceded by an asterisk (*) in the HTML report that the wizard generates.

Note: Within a specific table, columns are matched by **SASColumnName** and **SASColumnType** (numeric or character). Therefore, if a column's **SASColumnName** or **SASColumnType** has been changed by the user, the column might not be recognized as a supplied column. Also, if a supplied column needs to have its **SASColumnName** or **SAS ColumnType** changed in the supplied template table definition, then it is treated as a new column.

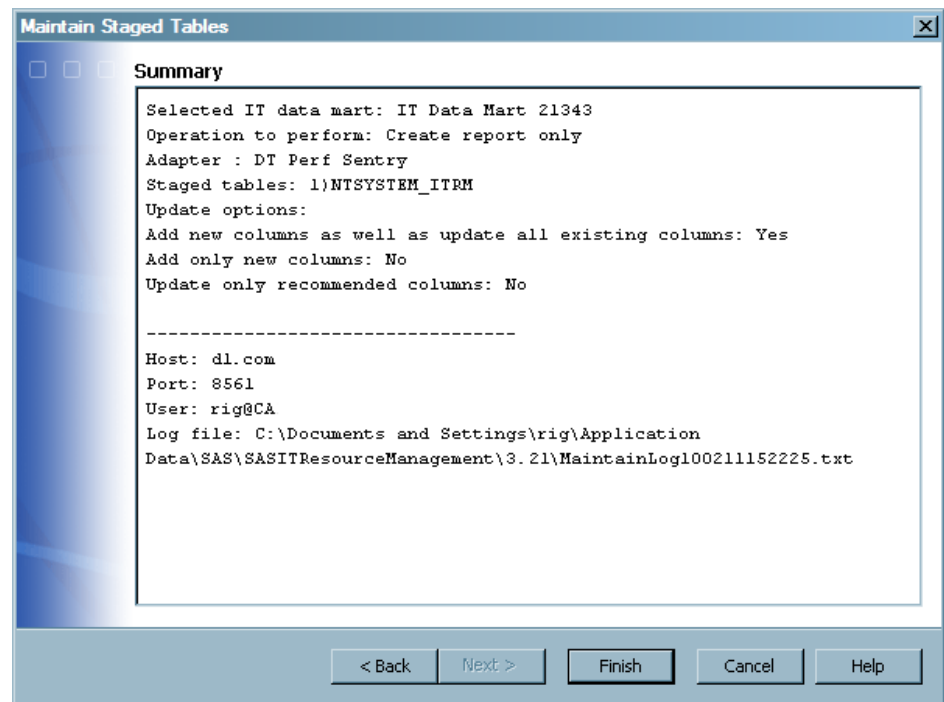
- b. If you choose to commit the updates that are detailed on this page, select one or more of the following **Update Options** that enable you to refine the scope of changes that you commit:
- **Add new columns as well as update existing columns** commits all of the changes that are detailed on this page of the wizard. This option must be deselected for the other options to be available.
 - **Add only new columns** commits only the new columns.
 - **Update only recommended columns** commits only the updates that affect the recommended columns in the associated template tables. For more information about recommended columns in template tables and determining which columns are recommended for a given template table, see [“Columns in Template Tables” on page 79](#).

Note: The **Update Options** are interdependent and the availability of each option in this window depends on whether the other options are selected or

deselected. For example, if **Add new columns as well as update existing columns** is selected, then the other options are not available. If this option is deselected, then the other options are available. Likewise, if **Add only new columns** or **Update only recommended columns** is selected, then **Add new columns as well as update existing columns** is not available for selection.

- c. Click **Next**.
7. View a summary of the impact analysis and update options that you selected for the staged tables.

Display 5.9 Summary Page in the Maintain Staged Tables Wizard



- a. Review the summary and ensure that it reflects the specifications and result that you expect.
- b. Click **Finish**. The wizard then generates an HTML report and saves it to your workstation.

Note: If there are no staged tables to maintain, the HTML report is not produced.

On the client, the report is located in a directory that is determined based on the APPDATA environment variable.

- On Windows XP, this location is typically **C:\Documents and Settings\userid\Application Data\SASHome\SASITResourceManagement\3.21\MaintainTableReportYYMMDDHHMMSS.htm**.
- On Windows 7, this location is typically **C:\Users\userid\AppData\Roaming\SASHome\SASITResourceManagement\3.21\MaintainTableReportYYMMDDHHMMSS.htm**.

If you selected **Commit updates with report** previously in the wizard, then the metadata for the selected staged tables is updated immediately. The physical data sets for these tables are then updated when the associated staging jobs are run.

Note: The data model version of the staged tables is updated and stamped only when you select both **Commit updates with report** and **Add new columns as well as update existing columns** in the wizard.

Using the Maintain Staged Tables Wizard to Update Staged Tables across All IT Data Marts

The Maintain Staged Tables wizard enables you to see whether any software changes to the template tables affect the staged tables that you use. You can also propagate these template table changes to their corresponding staged tables across all IT data marts at once. To do so, perform the following steps:

1. In the Folders tree view of SAS IT Resource Management, expand the following folders to locate the template table: **Products** ⇒ **SAS IT Resource Management** ⇒ **Adapters** ⇒ **<adapter name>** ⇒ **IT Template Tables** ⇒ **<domain category>** that corresponds to the template table for which you want to propagate updates to its corresponding staged tables.
2. Right-click the template table and select **Maintain Staged Tables** to open the wizard.
3. Complete the Maintain Staged Tables wizard. When you access the wizard from a specific template table as described in the previous steps, you are able to view and commit staged table updates across all IT data marts instead of within a single IT data mart. For more information about how to complete the Maintain Staged Tables wizard, see [“Using the Maintain Staged Tables Wizard to Update Staged Tables within an IT Data Mart” on page 85](#). This topic includes instructions for reading and completing each page of the wizard.

Maintaining Staged Tables in Batch Mode

About Maintaining Staged Tables in Batch Mode

In order to run the Maintain Staged Tables process in batch mode that you need to invoke the command from the server tier where the SAS IT Resource Management data tier has been installed. You cannot perform this action from a client install.

Note: You must use the Maintain Staged Tables wizard if you want to update selected staged tables that are affected by an update. The BatchMaintainTable command does not provide the ability to specify staged tables to update. For more information about maintaining specific staged tables, see [“Using the Maintain Staged Tables Wizard to Update Staged Tables within an IT Data Mart” on page 85](#).

Perform the following steps to maintain the tables in a batch process:

1. Open a command prompt and enter the following command:

```
cd C:\Program Files\SASHome\SASITResourceManagementDataTier\3.21
```

Note: The locations that are specified in this and in other sections of the following instructions assume that you have installed SAS IT Resource Management to a SASHOME location of **C:\Program Files\SAS**.

2. Execute the BatchMaintainTable command, specifying options as required.

TIP For best results, set all the options as needed but set the commit option to **N**. This setting produces a report that lists the actions that might be performed under the current settings of the command. However, if there are no staged tables to maintain, no HTML report is produced. After reviewing the report, if you are satisfied with the additions or updates that might be performed by using the

options that you specified, simply change the commit option to **Y** and rerun the command.

Syntax for the Batch Maintain Table Command Line

BatchMaintainTable -Commit Y | N

```
-AdapterName "adapter-name"
<-ITDataMartName "ITDataMartName">
<-TemplateName "TemplateName">
<-AddAndUpdateColumns Y | N>
<-AddNewColumnsOnly Y | N>
<-UpdateRecommendedColumnsOnly Y | N>
<-profile profile>
<-user userID>
<-password password>
<-host hostname>
<-port port>
<-domain domain>
<-?>
<-help>
```

Note: The ITDataMartName and TemplateTableName parameters are mutually exclusive.

To invoke the BatchMaintainTable command, use the following code:

```
"C:\Program Files\Java\jre1.5.0_12\bin\java.exe"
-Djava.system.class.loader=com.sas.app.AppClassLoader
-Dsas.ext.config="C:\Program Files\SASHome\SASPlatformObjectFramework
\9.3\sas.java.ext.config"
-Dsas.app.launch.config="C:\Program Files\SASHome
\SASITResourceManagementDataTier\3.21
\plugins\itrmdatatier\picklist"
-Dsas.app.repository.path="C:\Program Files\SASHome
\SASVersionedJarRepository\eclipse"
-Dlog4j.configuration="file:///C:\Program Files\SASHome
\SASDataIntegrationStudioServerData\4.3
\log4j.properties"

-Xmx1024m
-Xss1m
-classpath
"C:\Program Files\SASHome\SASVersionedJarRepository\
\eclipse\plugins
\sas.launcher.jar"
com.sas.solutions.itms.batch.maintainable.BatchMaintainTable
```

Note: If you installed SAS IT Resource Management on the C: drive in Windows, the preceding locations should be used. If you installed in another location, adjust the preceding locations accordingly.

BatchMaintainTable Options

Note: The ITDataMartName and TemplateTableName parameters are mutually exclusive. When the command is entered, make sure that each option is preceded in the command line by a hyphen.

- **Commit Y | N**
specifies whether to commit changes to the staged tables.

- AdapterName “*adapter-name*”
specifies the adapter whose staged tables you want to update.
Note: Enclose *adapter-name* in double quotation marks.
- ITDataMartName “*IT-data-mart-name*”
specifies the IT data mart where the adapter’s staged tables are stored. This option is required if the TemplateTableName option is not set.
Note: Enclose *IT-data-mart-name* in double quotation marks.
- TemplateTableName *template-table-name*
specifies the template table that you want to update. This option is required if the ITDataMartName option is not set.
- AddAndUpdateColumns Y | N
specifies whether to add new columns and update all existing columns. If the AddAndUpdateColumns option is N then either or both of the AddNewColumnsOnly and the UpdateRecommendedColumnsOnly options should be set to Y.
Note: Setting this option to Y is recommended.
- AddNewColumnsOnly Y | N
specifies whether to add only new columns. This option is required if the AddAndUpdateColumns option is not set to Y.
- UpdateRecommendedColumnsOnly Y | N
specifies whether to update only recommended columns. This option is required if the AddAndUpdateColumns option is not set to Y.
- profile *profile*
specifies the metadata server connection profile. This option can be used in place of the host, port, user, and password options.
- user *user-ID*
specifies the user login identity. This option is required if the profile option is not set or if the profile does not contain connection credentials.
- password *password*
specifies the user login password. This option is required if the profile option is not set or if the profile does not contain connection credentials.
- host *host-name*
specifies the metadata server host. This option is required if the profile option is not set.
- port *port-number*
specifies the metadata server port. This option is required if the profile option is not set.
- domain *domain*
specifies the user authentication domain.
- ? (or -help)
prints help information.

Note: By default, running the BatchMaintainTable command always produces a report of the results unless there are no staged tables to be maintained. In that case, no HTML report is produced.

Examples of the Batch Maintain Table Command Line

Example 1

This example shows how to run the Batch Maintain Table command on Windows. The command examines the staged tables in an IT data mart called “My IT Data Mart” that were generated for the BMC Performance Manager adapter. Any new or changed columns that apply to the staged tables for the BMC Performance Manager adapter are not performed, but are included in the report that is produced by the Batch Maintain Table command.

```
"C:\Program Files\Java\jre1.5.0_12\bin\java.exe"
-Djava.system.class.loader=com.sas.app.AppClassLoader
-Dsas.ext.config="C:\Program Files\SASHome\SASPlatformObjectFramework
\9.3\sas.java.ext.config"
-Dsas.app.launch.config="C:\Program Files\SASHome\SASITResourceManagementDataTier
\3.21\plugins\itrmdatatier\picklist"
-Dsas.app.repository.path="C:\Program Files\SASHome\SASVersionedJarRepository\eclipse"
-Dlog4j.configuration="file:///C:\Program Files\SASHome\
SASDataIntegrationStudioServerData\4.3\log4j.properties"
-Xmx1024m
-Xss1m
-classpath "C:\Program Files\SASHome\SASVersionedJarRepository\eclipse\plugins
\sas.launcher.jar"
com.sas.solutions.itms.batch.maintainable.BatchMaintainTable
-Commit N
-Adapter "BMC Perf Mgr"
-ITDataMartName "My IT Data Mart"
-AddAndUpdateColumns Y
-profile MyMetadataServerProfile
```

Example 2

This example shows how to run the Batch Maintain Table command in Windows. The command examines tables across all IT data marts on the application server that were generated for the BMC Performance Manager adapter and the PTCACHE template table. Only changes to the recommended columns are made. A report of these updates is produced.

```
"C:\Program Files\Java\jre1.5.0_12\bin\java.exe"
-Djava.system.class.loader=com.sas.app.AppClassLoader
-Dsas.ext.config="C:\Program Files\SASHome\SASPlatformObjectFramework
\9.3\sas.java.ext.config"
-Dsas.app.launch.config="C:\Program Files\SASHome\SASITResourceManagementDataTier
\3.21\plugins\itrmdatatier\picklist"
-Dsas.app.repository.path="C:\Program Files\SASHome\SASVersionedJarRepository\
-Dlog4j.configuration="file:///C:\Program Files\
SASHome\SASDataIntegrationStudioServerData\4.3\log4j.properties"
-Xmx1024m
-Xss1m
-classpath "C:\Program Files\SASHome\SASVersionedJarRepository\eclipse\plugins
```

```

\sas.launcher.jar"
com.sas.solutions.itms.batch.maintainable.BatchMaintainTable
-Commit Y
-Adapter "BMC Perf Mgr"
-TemplateTableName PTCACHE
-AddNewColumnsOnly Y
-UpdateRecommendedColumnsOnly Y
-profile MyMetadataServerProfile

```

Sample Batch Maintain Table Commands for Windows, UNIX and z/OS Operating Environments

Batch Maintain Table Command for Windows Operating Environments

```

"C:\Program Files\Java\jre1.5.0_12\bin\java.exe"
-Djava.system.class.loader=com.sas.app.AppClassLoader
-Dsas.ext.config="C:\Program Files\SASHome\SASPlatformObjectFramework
\9.3\sas.java.ext.config"
-Dsas.app.launch.config="C:\Program Files\SASHome\SASITResourceManagementDataTier
\3.21\plugins\itrmdatatier\picklist"
-Dsas.app.repository.path="C:\Program Files\SASHome\SASVersionedJarRepository
\eclipse"
-Dlog4j.configuration="file:///C:\Program Files\SAS
\SASDataIntegrationStudioServerData\4.3
\log4j.properties"

-Xmx1024m
-Xss1m
-classpath "C:\Program Files\SASHome\SASVersionedJarRepository\eclipse\plugins
\sas.launcher.jar"
com.sas.solutions.itms.batch.maintainable.BatchMaintainTable
-host <host-name>
-port (port-number)
-user <userID>
-password <password>
-domain DefaultAuth
-Commit N
-AdapterName "<adapter-name>"
-TemplateTableName <template-table-name>
-AddAndUpdateColumns Y
-AddNewColumnsOnly N
-UpdateRecommendedColumnsOnly N

```

Note: Double quotation marks are needed in this command line wherever a part of an address (for example, “Program Files”) contains a space.

Batch Maintain Table Command for UNIX Operating Environments

This example is enclosed in a shell script. It was run on the H64 UNIX operating environment.

```

#!/bin/sh

export SASHOME=<full-path-to-SASHOME directory>

```

```

echo $SASHOME is $SASHOME
export DISPLAY=<hostname:displaynumber.screennumber>
echo DISPLAY is $DISPLAY

cd $SASHOME/SASITResourceManagementDataTier/3.21

$JAVA_HOME/jre1.5.0.08/bin/java \
-Djava.system.class.loader=com.sas.app.AppClassLoader \
-Dsas.ext.config=$SASHOME/SASPlatformObjectFramework/9.3/sas.java.ext.config \
-Dsas.app.launch.config=$SASHOME/SASITResourceManagementDataTier
/3.21/plugins/itrmdatatier/picklist \
-Dsas.app.repository.path=$SASHOME/SASVersionedJarRepository/eclipse \
-Dlog4j.configuration="file://$SASHOME/SASITResourceManagementDataTier
/3.21/log4j.properties" \
-Djava.awt.headless=false \
-Xmx1024m \
-Xss1m \
-classpath $SASHOME/SASVersionedJarRepository/eclipse
/plugins/sas.launcher.jar \
com.sas.solutions.itms.batch.maintaintable.BatchMaintainTable \
-host <metadata-server-hostname> \
-port <metadata-server-port-number> \
-user <unrestricted-metadata-server-userID> \
-password <password-for-specified-user> \
-domain DefaultAuth \
-Commit N \
-AdapterName "BMC Perf Mgr" \
-TemplateTableName PTCACHE \
-AddAndUpdateColumns Y \
-AddNewColumnsOnly N \
-UpdateRecommendedColumnsOnly N
    
```

Batch Maintain Table Command for z/OS Operating Environments

This example is enclosed in a shell script. It was run on the z/OS operating environment.

```

#!/bin/sh

echo JAVA_HOME is $JAVA_HOME
export SASHOME=<full-path-to-SASHOME-directory>
echo $SASHOME is $SASHOME
export DISPLAY=<hostname:displaynumber.screennumber>

echo DISPLAY is $DISPLAY

cd $SASHOME/SASITResourceManagementDataTier/3.21
$JAVA_HOME/bin/java \

-Djava.awt.headless=false \
-Djava.system.class.loader=com.sas.app.AppClassLoader \
-Dsas.ext.config=$SASHOME/SASIntelligencePlatformLogicalTypes
/9.3/sas.java.ext.config \
-Dsas.app.launch.config=$SASHOME/SASITResourceManagementDataTier
/3.21/plugins/itrmdatatier/picklist \
-Dsas.app.repository.path=$SASHOME/SASVersionedJarRepository/eclipse \
-Dlog4j.configuration="file://$SASHOME/SASITResourceManagementDataTier
/3.21/log4j.properties" \
    
```

```

-Xmx1024m \
-Xss1m \
-classpath $SASHOME/SASVersionedJarRepository/eclipse/plugins/sas.launcher.jar \
com.sas.solutions.itms.batch.maintainable.BatchMaintainTable \
-host <metadata-server-hostname> \
-port <metadata-server-port-number> \
-user <unrestricted-metadata-server-userID> \
-password <password-for-specified-user> \
-domain DefaultAuth \
-Commit N \
-AdapterName "BMC Perf Mgr" \
-TemplateTableName PTCACHE \
-AddAndUpdateColumns Y \
-AddNewColumnsOnly N \
-UpdateRecommendedColumnsOnly N

```

Chapter 6

Formulas

About Formulas	97
What Is a Formula?	97
How Can a Formula Be Used?	97
How Is a Formula Created?	98
How Can Formulas Be Accessed?	98
How to Access the Expression Builder Window	99
Working with Formulas	100
Creating a Formula	100
Delete a Formula	102
Modify a Formula	102
Consequences of Changing a Formula	103
Rename a Formula	104

About Formulas

What Is a Formula?

A formula is an expression that calculates the value for a computed column in a staged or aggregation table.

There are two types of formulas: supplied formulas and user-defined formulas.

- Supplied formulas are provided by SAS IT Resource Management and are loaded along with the SAS IT Resource Management software at installation time. Supplied formulas cannot be modified. However, you can copy a supplied formula and modify the copied version, thus creating a user-defined formula.
- User-defined formulas are created by the data administrator or the performance analyst.

How Can a Formula Be Used?

Formulas can be used in all SAS IT Resource Management tables in order to calculate new computed columns.

Computed columns can be used for various purposes, such as to standardize the measurement units or time zones of related data or to establish datetime-related constructs such as WEEKDATE and SHIFT.

- Computed columns are supplied with the template tables for supplied adapters.
- Computed columns can be generated and added to staged or aggregation tables.


For more information about computed columns, see [“Specify Computed Columns” on page 153](#).

How Is a Formula Created?

There are two ways to create a user-defined formula.

- A user-defined formula can be created by using the New Formula wizard, which is invoked by selecting **Formula** from the **File** ⇒ **New** ⇒ **IT Resource Management** menu of SAS IT Resource Management.
- A user-defined formula can be created by copying a supplied formula and pasting it into a folder to which you have Write privileges. The newly copied formula will retain its name unless it is pasted into a folder that already contains a formula with the same name. In that situation, the formula will be named "Copy of **formula name**." After it is pasted into the target folder, the formula can be renamed and modified.

How Can Formulas Be Accessed?

Formulas () can be accessed from several folders in the Folders tree of SAS Data Integration Studio and can be used when specifying the expression that defines a computed column.

- **Products** is a folder that contains formulas that are supplied with SAS IT Resource Management software. These supplied formulas cannot be changed. However, they can be copied to another folder where they can be modified.

To access these supplied formulas from the Folders tree, expand **Products** ⇒ **SAS IT Resource Management** ⇒ **IT Formulas**.

- **Shared Data** is a folder that contains formulas that are supplied with SAS IT Resource Management software, as well as those user-defined formulas that can be shared by other users at your site. Formulas that are in this folder can be modified.

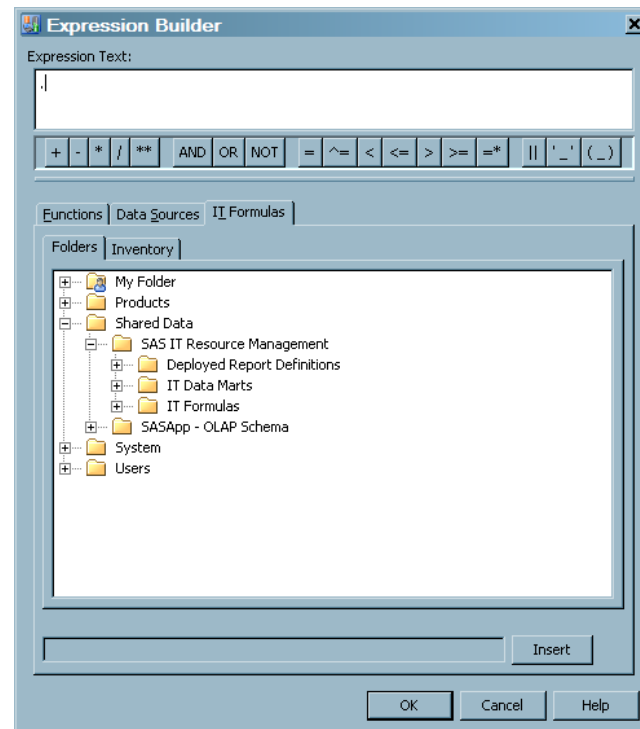
To access these formulas from the Folders tree, expand **Shared Data** ⇒ **SAS IT Resource Management** ⇒ **IT Formulas**.

- **My Folder** is a folder that contains user-defined formulas that you have saved in that location. **My Folder** is the private folder of the person, or owner, who is currently logged on to SAS IT Resource Management. The contents, including formulas, are not accessible to anyone other than the owner.
- The **Users** folder contains the private folders of all the users who are registered users of the metadata repository to which you are connected. The only folder that you can access in the **Users** folder is your own, private folder. It is the folder that belongs to the person who is logged in to SAS IT Resource Management.

The following display shows the **IT Formulas** tab of the Expression Builder window. You can expand **My Folder**, **Shared Data**, or your private folder within the **Users** folder and use any of those formulas in your expression. Alternatively, you can select the

Inventory tab and from the list that is displayed, use any of those formulas in your expression.

Display 6.1 IT Formulas Tab of Expression Builder



Note: The **IT Formulas** tab of the Expression Builder window is available only when you are editing columns of a staged or aggregation table from the Properties dialog box in the process flow diagram.

How to Access the Expression Builder Window

You can use the Expression Builder window to add a new computed column to a table or to modify an existing computed column in a table.

To access the Expression Builder window in order to add a new computed column to a table, perform the following steps:

1. From the IT data mart that contains the table, double-click the job that uses the table. The job opens on the **Diagram** tab of the Job Editor window.
2. Select the table to which you want to add a column. Right-click the table to display the Properties dialog box of that table and click **Properties**.
3. Click the **Columns** tab to display the columns that make up the table.
4. Click **New** to add a new column to the table. From the drop-down menu, click **Computed column**.
5. Double-click the enabled **Expression** field for the new field to display an ellipsis (...).
6. Click the ellipsis (...). The Expression Builder window appears.

To access the Expression Builder window in order to modify a computed column in a table, perform the following steps:

1. From the IT data mart where the table is stored, double-click the job that uses the table. The job opens on the **Diagram** tab of the Job Editor window.
2. Select the table whose column you want to modify. Right-click the table to display the Properties dialog box of that table and click **Properties**.
3. Select the **Columns** tab to display the columns that make up the table.
4. Select the **Expression** field of the column that you want to change.
5. Double-click the enabled **Expression** field for the field to display an ellipsis (...).
6. Click the ellipsis. The Expression Builder window appears.

For more information about the Expression Builder window, press the F1 key or select **Help** from within that window.

Working with Formulas

You can create, rename, delete a formula. You can also modify a formula and change its expression.

Creating a Formula

How to Create a Formula

You can create a formula in two ways:

- by invoking the New Formula wizard
- by copying an existing formula to a folder, renaming it, and modifying it

You can create a formula and save it in any of the **Formulas** folders to which you have Write privileges. For example, you can create a formula in the **Formulas** folder of **Shared Data** or in your private folder within **My Folder**.

If you want to use a formula in a wizard, you can specify the location of that formula. If the formula is not found in your specified location, then the wizard will search for the formula in this location: **Shared Data/SAS IT Resource Management/IT Formulas**. If the formula is still not found, the wizard will search for it in this location: **Products/SAS IT Resource Management/IT Formula**.

Create a Formula by Using the New Formula Wizard

To create a formula using the New Formula wizard, perform the following steps:

1. From the SAS Data Integration menu bar, select **File** ⇒ **New** ⇒ **IT Resource Management** ⇒ **Formula** to open the New Formula wizard.
2. On the General Information page of the New Formula wizard, specify the name, description, and location of the new formula.
 - The **Name** field can contain a maximum of 60 characters. A value is required in this field and it must be unique within its folder.

Note: The wizard automatically generates a name for this new formula. You can either delete the system-generated name and enter a more meaningful name for this formula or you can retain the system-generated name.

- The **Description** text box can contain no more than 200 character. A value is optional in this field.
- The **Location** field is automatically generated and, by default, specifies / **Shared Data/SAS IT Resource Management/IT Formulas**. Click **Browse** to navigate to another folder to use as the location for your new formula.

TIP User-defined formulas are more easily accessed if they are all stored in a location that is identified as a “Formulas” folder.

3. On the Expression page of the New Formula wizard, enter the expression that specifies the combination of functions and mathematical operations that are used to derive a value. A value is required in this field.

You can enter the expression for a formula in any one of these three ways:

- Type the expression in the text area of the **Expression** field.
- Copy and paste the expression text from the **Expression** tab of an existing formula to the **Expression** text area of the new formula.
- Copy and paste the computation from the **Expression** tab of the computed column into the **Expression** text area of the new formula.

In addition to the conventional rValue expression (where the expression consists of code that is appropriate only for the right-hand side of an assignment statement), SAS IT Resource Management also supports more complex expressions. For example, your expression can use SAS code that might include loops, IF statements, and so on. As shown in the following display, this code must be written in valid SAS DATA step syntax.

Display 6.2 Expression Page of the New Formula Wizard

4. The final page of the New Formula wizard is a summary page that displays the following information about the formula that you specified. Click **Finish** to store the formula that you created in the folder that you specified in the **Location** field.

Create a Formula by Copying an Existing Formula

To create a formula by copying an existing formula, perform the following steps.

1. In the Folders tree of SAS IT Resource Management, navigate to an **IT Formulas** folder (or to a folder that contains formulas) to locate the formula that you want to use as a template for the new formula.

TIP User-defined formulas are more easily accessed if they are all stored in a location that is identified as a “Formulas” folder.

2. Right-click the formula that you want to copy. From the list that is displayed, select **Copy**. Then navigate to the folder where you want to store your new formula and click **Paste**.

If you are copying and pasting the formula to the same folder, then the name of the new formula is retained unless a formula of the same name already exists at this location. In that situation, **Copy of** is appended to the original name of the formula. This ensures that the formula name is unique within its folder.

Note: You can select the **Copy to Folder** option, which opens the Select a Location dialog box in which you can navigate to the folder where you want to store the copied formula. You can also select the **Move to Folder** option, which opens the Select a Location dialog box in which you can navigate to the folder where you want to store the formula.

3. On the new copy of the formula, modify the fields as required. For information about changing the fields of a formula, see [“Modify a Formula” on page 102](#).

Delete a Formula

You can delete formulas only from the folders for which you have Write privileges.

To delete a user-defined formula, perform the following steps:

1. In the Folders tree of SAS data Integration Studio, navigate to an **IT Formulas** folder or a folder where the formula that you want to delete is located.
2. Right-click the formula. From the list that is displayed, select **Delete**.
3. In the confirmation dialog box, click **Yes** to delete the formula. Otherwise, click **No**.

Deleting a formula does not change the expression that used the formula, and the columns that are computed using that expression are not affected. The columns will continue to be created according to the specified expression. However, when a formula that is used in an expression no longer exists, then modifications to that expression can no longer be performed simply by changing the formula. The modifications must be changed in every column where the expression is used.

Modify a Formula

You can change any formula to which you have Write privileges. To do so, perform the following steps:

1. Double-click the formula to open the Properties dialog box, where you can change the name, description, location, and expression of the formula.
 - Select the **General** tab to change the name, description, or location of the formula.
 - To change the name or description of the formula, type your modifications in the corresponding fields.
 - To change the location of the formula, click **Browse** to open the Select a Location dialog box where you can navigate to another folder.
2. Select the **Expression** tab to change the combination of functions and mathematical operations that are used to derive a value.

If you modify the expression text for a formula, then the system checks to see whether the original formula is used by any computed column of any table. If the formula is being used, an Update Formula dialog box opens that lists the computed columns that will be affected by the change. You can select **Yes** to change the formula in the computed columns that use it or **No**. If you click **Yes**, then the modified expression for the formula is saved and applied in all instances where it is used.

Note: Computed columns in staged tables or aggregation tables will be calculated using the latest version of the formula.

3. Select the **Notes** tab to enter additional information about the formula.

CAUTION:

The Advanced tab contains information about the metadata for this formula. The Authorization tab contains information about the groups, users, and permissions that are in effect for this formula. Do not change the information that is displayed on these tabs.

4. Click **OK** to make your changes take effect.

Consequences of Changing a Formula

Formulas can be shared across multiple computed columns and across multiple IT data marts. For example, the SHIFT computed columns for a particular IT data mart can all be based on the same SHIFT formula. If you want to change the site-specific definition of SHIFT, you can make a change to the SHIFT formula and click **OK**. If this expression is used in any other formula, the Update Formula dialog box appears and displays the computed columns where this formula is used. Click **Yes** to cause all computed columns that use the formula to inherit the change. Otherwise, click **No**.

Computed columns in staged or aggregation tables are calculated using the latest version of the formula. In other words, if you change the expression for a computed column in a staged table or an aggregation table and redeploy the job, then the next time the job is run, the values for that column in that staged table or aggregation table will be computed according to the new expression. This computation occurs whether the expression is from a formula or only for that computed column.

CAUTION:

If the expression of a computed column is changed and if that column was already used as input to a subsequent transformation that in turn created a target table, then the values in that target table are not automatically recreated.

To make the old and new values of the computed columns consistent with each other, redeploy and rerun all the previous ETL jobs that contain the new formula, according to the following guidelines.

- Case 1: The change to the expression of a formula should be reflected in all the tables, including the aggregation tables. This situation might occur if an erroneous expression was entered and needs to be corrected. Previously aggregated data is affected and must be changed.

Solution: Purge the aggregation tables. Then redeploy and rerun all the staging jobs and all the aggregation jobs that contain the data that uses the altered expression to create a new column. Thus, you recreate the aggregation tables.

- Case 2: The change to the expression of a formula should be reflected only from this time forward. This situation might occur if there has been a change to the billing rate. Previous aggregated data is not affected.

Solution: Redeploy the staging job that uses the altered expression. You do not have to redeploy and rerun the aggregation jobs whose source tables use the altered expression to create a new column.

For example, consider the following situation:

- There is a computed column RATE in a staged table, where $\text{RATE} = \text{BYTES} / \text{SECOND}$.
- There is an aggregation with statistics based on that RATE column (for example, the mean RATE).
- Data for the RATE and the mean RATE columns has already been aggregated.

If you subsequently change the expression for RATE to $\text{RATE} = \text{BYTES} / \text{MINUTE}$, and redeploy the job, in the next ETL, the values of the staged table's RATE column will be computed according to BYTES/MINUTE. However, the values in the aggregation table that already exist (that is, they were already aggregated) will still reflect the previous BYTES/SECOND calculation. According to the guidelines in Case 1, if you want the statistics that are generated to reflect the rate change in the aggregated data, you should rerun all the associated staging and the aggregation jobs.

Rename a Formula

Renaming a formula does not affect the tables that use it. To rename a formula, perform the following steps:

1. In the Folders tree of SAS data Integration Studio, navigate to the **IT Formulas** folder where the formula that you want to rename is located.
2. Right-click the formula. From the list that is displayed, select **Rename**. The name of the formula is highlighted.
3. Enter the new name of the formula. The name can contain no more than 60 characters. The name must be unique within its folder.
4. Click **Enter**.

Chapter 7

Staging the Data

About Staging the Data	106
Why Stage the Data?	106
Methods for Staging the Data	106
Working with Staging Transformations	107
What is a Staging Transformation?	107
Properties of Staging Transformations	108
Add a Staging Transformation to a Staging Job	109
Edit a Staging Transformation	112
Delete a Staging Transformation	112
Working with Staged Tables	112
What is a Staged Table?	112
Properties of a Staged Table	113
Task List for Staged Tables	114
Create Staged Tables	114
Add an Existing Staged Table to a Job	120
Erase a Staged Table	120
Delete a Staged Table from a Tree View	121
Delete a Staged Table from a Job	121
Purge the Content of a Staged Table	122
Modify a Staged Table	122
Columns in Staged Tables	124
Computed Columns in Staged Tables	126
Staging VMware Data Using Local Time	126
About Staging VMware Data Using Local Time	126
Change the Formula Expression for the IT Formula VMwareDatetime	127
Change the Expressions for Other Datetime-related	
Computed Columns in Your Staged Tables	127
Understanding External Names of Staged Tables and Columns	127
About External Names of Staged Tables and Columns	127
Change the External Name of a Table or a Variable	128
View the External Names of Tables	128
View the External Names of Columns	129
External Names for MXG Adapters	129
Insufficient Disk Space	129

About Staging the Data

Why Stage the Data?

IT performance data is created and processed by unique IT systems, methods, management utilities, and code. Therefore, SAS IT Resource Management must stage the raw data that is created by these various utilities so that the IT performance data is standardized and comparable. For example, the formats for timestamps and units of measurement must be standardized, duplicate data must be processed based on business needs, and so on.

SAS IT Resource Management uses the following two components of a job to prepare and stage raw data:

- staging transformations (including user-written staging transformations)
- staged tables

Note: The VMware vCenter adapter uses two staging transformations (VMware Data Acquisition and VMware vCenter) to prepare and stage raw data. For information about the VMware Data Acquisition staging transformation, see [“VMware vCenter Data Acquisition Jobs” on page 351](#).

Staging transformations generate the code that extracts raw data, transforms it to a standardized format, and loads it into staged tables. These staged tables are based on template tables that are supplied for supported adapters. For more information about template tables, see [“Working with Template Tables” on page 75](#).

You can review the specifications of staging transformations and staged tables and can configure elements of the extraction and staging process to accommodate the requirements of your raw data and business environment. For more information about staging transformations, see [“What is a Staging Transformation?” on page 107](#). For more information about staged tables, see [“What is a Staged Table?” on page 112](#).

Methods for Staging the Data

To administer the staging process and specify the parameters for your environment, you must create and deploy one or more jobs that include an associated staging transformation for the adapter that you use. Staging transformations are included in staging jobs in order to create the SAS programs that stage the data and populate the staged tables. These staging jobs create SAS staged tables in the IT data mart. The SAS staged tables serve as the basis for all performance management, capacity planning data management, analysis, and reporting that you perform when using SAS IT Resource Management.

SAS IT Resource Management provides the following methods for staging IT performance data to create and populate staged tables:

- using the Adapter Setup wizard to create staging transformations and staging jobs automatically for a given adapter. When these staging transformations are executed within a job, a set of staged tables is created based on the parameters set in the Adapter Setup wizard. For more information about using the Adapter Setup wizard to stage data, see [“Using the Adapter Setup Wizard” on page 179](#).

- creating and modifying jobs with staging transformations manually and using the Staged Table wizard to create customized staged tables. For more information about performing these tasks, see the following topics:
 - [“Add a Staging Transformation to a Staging Job” on page 109](#)
 - [“Create Staged Tables” on page 114](#)

You can also create staging transformations and staged tables to stage data from any input source based on the unique specifications of any business environment. SAS IT Resource Management provides a user-written staging transformation for user-written code. For more information about the user-written staging transformation, see [“User-Written Staging Transformations” on page 233](#).

SAS Data Integration Studio tools such as the Source Designer and Process Designer can also help create the staging code and the staging job that is necessary to support any input source. For more information about using these tools to create user-written staging code, see [“User-Written Staging Transformations” on page 233](#).

Working with Staging Transformations

What is a Staging Transformation?

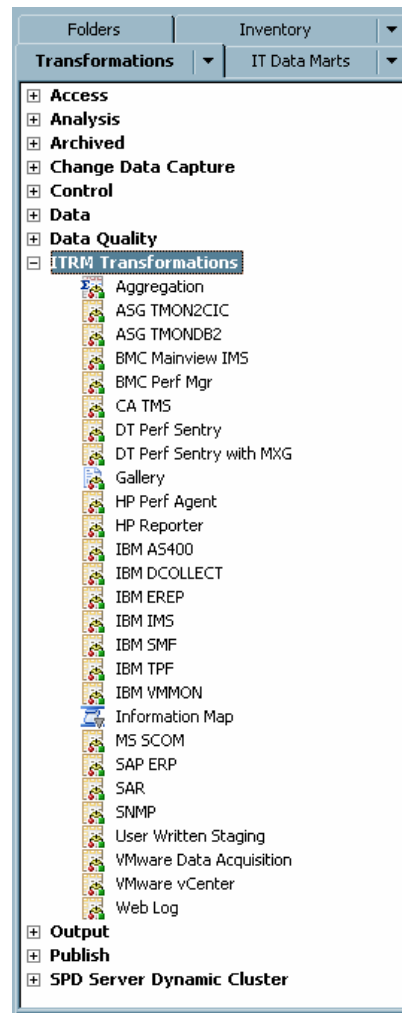
A staging transformation generates a SAS program that loads raw data into staged tables. Before loading raw performance data into staged tables, the staging transformation performs the processing that is necessary to create a SAS program that converts the raw data into standardized, normalized values. For example, the staging transformation might perform the following functions:

- provide code that supports the input of single files, directories, or databases
- detect and handle duplicate data
- normalize measurement units
- standardize timestamps
- generate computed columns from the input data

When a staging transformation is executed by a job in SAS IT Resource Management, the SAS program accesses the parameters that are specified in the staging transformation and the metadata for the staged table that it will populate. The SAS program uses this information to populate the staged table that can be used as input to an Aggregation transformation or other transformations. Every time a staging transformation is executed, any existing data in the staged table is removed, and the new data is loaded.

You can also modify a staging transformation to include additional staged tables that are supported by a given adapter. This feature enables the system to process data from various staged tables for a given data source via the same staging transformations and aggregations. For more information about aggregations, see [“Working with Aggregation Transformations” on page 170](#).

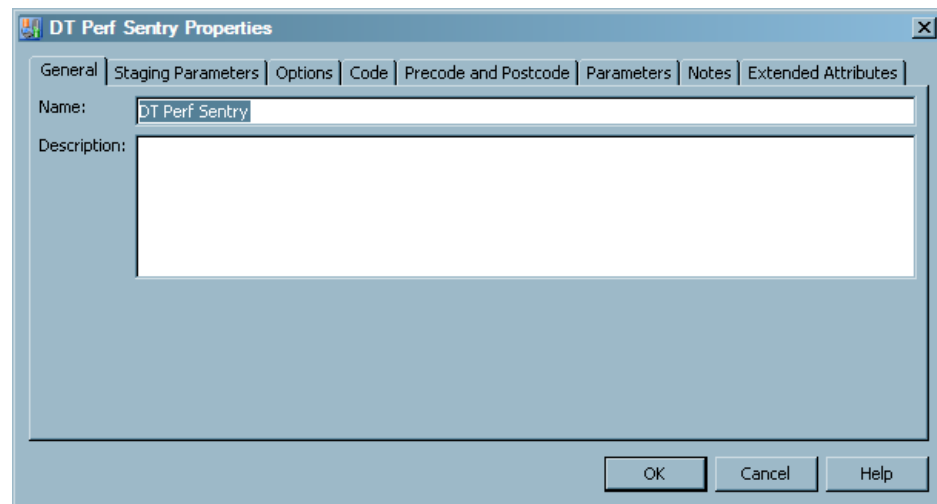
SAS IT Resource Management provides a unique staging transformation for each adapter that it supports. The staging transformations that are supplied for each adapter are available in the **ITRM Transformations** folder of the Transformations tree view of SAS IT Resource Management, as shown in the following display.

Display 7.1 Supplied Staging Transformations in ITRM Transformations Folder

Properties of Staging Transformations

Staging transformations specify the staging parameters that are pertinent for the adapter on which they are based. Because raw data sources differ in format, file structure, and type of data collected, the staging parameters that are necessary for each adapter vary as well. Staging parameters enable you to indicate how you want to handle certain types of data based on your needs. For example, you can specify how to handle duplicate or future data in a raw data source. Staging parameters can also have different default values based on the adapter.

To access the staging transformation properties, right-click a staging transformation in the process flow diagram and select **Properties**. A Properties dialog box appears and enables you to view the staging transformation properties and specifications for the staging code that the transformation generates.

Display 7.2 Sample Properties Dialog Box for a DT Perf Sentry Staging Transformation

The Properties dialog box contains the following tabs:

- The **General** tab displays information that identifies the staging transformation such as name and description.
- The **Staging Parameters** tab displays the staging parameters that are pertinent for the adapter on which the staging transformation is based. The parameters that are available on this tab vary based on the requirements of each adapter. You can click the value field of a parameter to modify the value. If there is a defined set of values available for this parameter, a drop-down list appears and enables you to select a value. For more information about staging parameters, see “[Staging Parameters](#)” on [page 356](#).

Note: If you modify the value of any parameter in the source code of a job, that modification is not updated in the SAS Metadata Repository, and the new value is not reflected on the **Staging Parameters** tab. However, you can save the modified version of the source code to the local file system.

- The **Options** tab displays information about how to further customize and generate the code.
- The **Code** tab enables you to manage the code that is generated for the transformation.
- The **Precode and Postcode** tab enables you to specify that user-written code should be inserted at the beginning or end of the current job or transformation.
- The **Parameters** tab enables you to manage prompts or prompt groups.
- The **Notes** tab displays any notes or documents that are associated with the transformation.
- The **Extended Attributes** tab displays the custom properties that are available for the transformation.

For more information about the tabs that are available from the Properties dialog box, see the Help for SAS IT Resource Management.

Add a Staging Transformation to a Staging Job

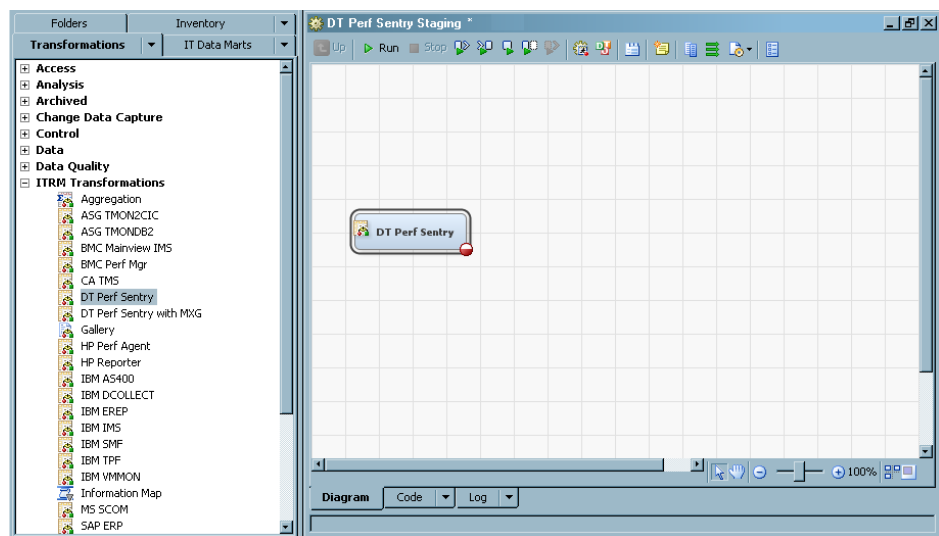
You can use the Adapter Setup wizard to create staging jobs automatically for a supplied adapter. The Adapter Setup wizard provides a convenient way to create staging jobs as

well as other jobs and the components that are necessary to stage and summarize the IT resource data that an adapter loads. The wizard guides you through the process of specifying the staging parameters, and it creates the necessary staging transformation, staging job, and other jobs that are required for the adapter that you select. For more information about using the Adapter Setup wizard, see [“About the Adapter Setup Wizard” on page 177](#).

You can also create a staging job manually by adding a staging transformation to a job. To do so, perform the following steps:

1. Open the job that is to contain the new staging transformation. The job shows on the **Diagram** tab of the Job Editor window.
- Note:* If you want to add a staging transformation to a new job, you must first create the job. For more information about how to create a job, see [“Create a Job” on page 268](#).
2. In the Transformations tree of SAS IT Resource Management, expand the **ITRM Transformations** folder and locate the staging transformation for the adapter that you choose.
 3. Drag and drop the staging transformation onto the Job Editor window. The staging transformation appears in the process flow diagram for the job.

Display 7.3 Sample Job with a New Staging Transformation



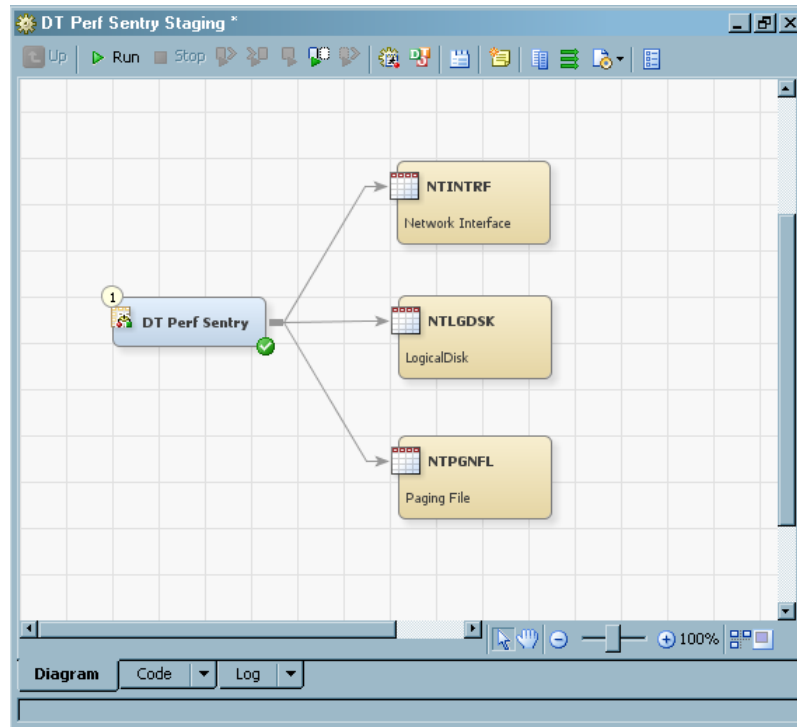
4. Right-click the staging transformation and select **Add New Staged Table** to open the Staged Table wizard. This wizard enables you to specify the parameters for the staged tables that the staging transformation code populates. For more information about using the Staged Table wizard, see [“Create Staged Tables” on page 114](#).

Note: A staging transformation can process output for multiple staged tables. To add more staged tables, repeat this step as needed.

5. Right-click the staging transformation and select **Properties**.
6. Use the **Staging Parameters** tab to specify the staging parameters for your data such as the location of the raw data, how to handle duplicate data, how to handle future data, and so on. You must specify all parameters that are marked with an asterisk (*) before the staging transformation is complete. When the transformation is complete and ready to run in a staging job, the staging transformation object is marked with a green check in the process flow diagram.

For more information about the staging parameters for each adapter, see “[Staging Parameters](#)” on page 356.

Display 7.4 Sample Job with a Complete Staging Transformation and Staged Tables



Note: The location of the raw data input for the staging transformation is specified in the staging transformation properties. Thus, a staging transformation object does not have an input object in the process flow diagram (with the exception of the VMware vCenter staging transformation that has the data acquisition table as input).

7. If you are adding an ASG TMON2CIC or IBM SMF staging transformation, you must specify the associated spin library. Use the **Spin Library** tab (available only for these two staging transformations) to view and modify the parameters for the corresponding spin library. For more information about setting the spin library parameters, see “[Data Sources Supported by SAS IT Resource Management Adapters](#)” on page 290.

8. Run the staging job to generate the staging code, execute the code, and load the staged tables.

Note: You can also generate the staging code without executing it. To do so, click the **Code** tab of the Job Editor window. You can then select from the **Code generation mode** options on this tab to edit the code manually before executing it.

9. Click the **Log** tab of the Process Designer window to check the SAS log and confirm that there were no errors or warnings during processing.

After the staging job executes successfully, you can view the staged data in the staged tables. To do so, right-click a staged table in the process flow diagram and select **Open**.

Edit a Staging Transformation

To edit the parameters that a staging transformation uses to locate and stage raw data, perform the following steps:

1. In the IT Data Marts tree, double-click the job that contains the staging transformation that you want to edit.
2. Right-click the staging transformation in the process flow diagram and select **Properties**. A Properties dialog box displays and enables you to modify the various parameters that are specified for the staging transformation and the staging code that it generates.

Note: The name of a staging transformation should not contain double quotation marks.

For more information about staging transformation properties, see [“Properties of Staging Transformations” on page 108](#).

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

Delete a Staging Transformation

Deleting a staging transformation disassociates it permanently from the staged tables in the job.

CAUTION:

If you delete a staging transformation and then try to re-create or reconnect it to the staged tables, the staged tables might lose computed columns or experience other errors. For best results, make a copy of the staging job before deleting or modifying the staging transformation or any other components that you are not certain that you want to lose.

To delete a staging transformation from a job, perform the following steps:

1. In the IT Data Marts tree, double-click the job that contains the staging transformation that you want to delete.
2. Right-click the staging transformation in the process flow diagram and select **Delete**. The staging transformation object is deleted from the job. However, the staged tables that were associated with the staging transformation remain in the job.
3. Select **File** ⇒ **Save** to save your changes, or close the job and select **Yes** when prompted to save your changes.

Working with Staged Tables

What is a Staged Table?

A staged table contains data that has been extracted from an input data store by a staging transformation and has been rendered into a SAS data set that is suitable for further transformation. The structure and metadata for most staged tables that SAS IT Resource Management provides are based on corresponding template tables for specific adapters

and domain categories. For more information about template tables, see [“Working with Template Tables” on page 75](#).

The data in a staged table is accumulated according to rules that are specified by a staging transformation. The IT performance data resides in the staged table after the data is refined and loaded into a SAS data set that is ready for aggregation and reporting. Likewise, the staged table determines how the data appears after it has completed the staging process.

You can create a staged table by using the Adapter Setup wizard, the Staged Table wizard, or a user-written staging transformation. When a staged table is created, it is associated with a staging transformation in a staging job and stored in the same subfolder of the IT Data Marts tree as the staging job by default. This enables you to keep track of the staged tables and their corresponding jobs as they are stored together.

Note: If you manually create your own user-defined tables using the Source Designer or Target Designer, then consider saving these tables in a similar manner to maintain consistency.

When a staging job executes, the staging transformation reads the raw IT performance data from an input data source, processes it, and loads it into the corresponding staged tables. These tables can then be input into Aggregation transformations.

The staged table is generally used as input to an Aggregation transformation. However, a user can change this process flow by using the manual methods that are available in SAS IT Resource Management.

SAS IT Resource Management supplies adapters that provide staging transformations, template tables, and staged tables for many diverse data sources.

Note: The New Table wizard of SAS Data Integration Studio can create other tables that stage raw data from any other input data sources.

For information about the adapters that SAS IT Resource Management supports, see [“Supported Adapters” on page 68](#).

Properties of a Staged Table

The metadata about a staged table, such as how raw data is computed and formed into a SAS data set structure, is available from the Properties dialog box of a staged table.

To view the properties of a staged table, complete the following steps:

1. Navigate to the folder in the IT Data Marts tree that contains the staging job and staged table.
2. Open the staging job that includes the staged table.
3. Right-click the staged table in the process flow diagram and select **Properties**.

Note: You can also right-click the staged table in the tree view of SAS IT Resource Management to access the staged table properties. However, the set of properties that are available when using this method differs slightly from the properties that are available from the process flow diagram. For best results when viewing and modifying relevant properties, access the properties from the staged table object in the process flow diagram as explained in the preceding steps.

The following list includes all of the staged table properties and indicates whether a property is available only from the staged table in the tree view or the staged table in the process flow diagram:

- The **General** tab displays information that identifies the staged table name, description, location, and associated responsibilities.
- The **Columns** tab displays the metadata for each column in the staged table. For more information about the properties that are available from the **Columns** tab of a staged table, see “Columns in Staged Tables” on page 124.
- The **Indexes** tab displays any indexes that have been created for this staged table. For each column, the name, description, length, and type are stored.
- The **Keys** tab displays any keys that have been created for this staged table. However, the values displayed on this tab are ignored by SAS IT Resource Management.
- The **Parameters** tab displays any prompts or groups that have been created for this staged table. This tab is available from the tree view only.
- The **Physical Storage** tab displays information about the physical name of the table and the library and DBMS where it is located.
- The **Options** tab displays table options such as compressing observations in a SAS data set, additional security options, and rules for overwriting data sets.
- The **Notes** tab displays any notes or documents that are associated with this staged table.
- The **Extended Attributes** tab displays information about any custom property that is not part of the standard metadata for this staged table. This tab includes any external names for a table, if available.
- The **Authorization** tab, available only from the tree view, displays the settings that define access to the staged table. This tab is available from the tree view only.
- The **ITRM Options** tab displays the adapter, template table, data model version, and predecessor tables that are associated with the staged table. This tab is available from the process flow diagram only.

For information about the contents of these tabs, click **Help** in the dialog box.

Task List for Staged Tables

When you right-click a staged table, you gain access to several actions that you can perform for that table such as analyze, update metadata, delete, copy, move, export, and so on. The tasks that are available differ depending on whether you right-click the staged table in the tree view or the process flow diagram. For information about these tasks, see the Help that is available for these topics in SAS IT Resource Management.

Create Staged Tables

SAS IT Resource Management does not enable you to create a staged table without associating it to a staging transformation. Thus, you must have a staging job that includes a staging transformation before creating a staged table.

If you want to create staged tables for the BMC Perf Mgr, DT perf Sentry, or SNMP adapters, other than the tables that are included in the staged table templates for those adapters, you can specify and then run macros to do so. These macros create the

metadata for the staged table and attach that metadata to the staging transformation in your job.

- To create staged tables for BMC Perf Mgr, see “%RMMSBMCP” on page 519.
- To create staged tables for DT perf Sentry, see “%RMMSDTPS” on page 521.
- To create staged tables for SNMP, see “%RMSSSNMP” on page 523.

Note: For best results, always create or add staged tables to a job from the Staged Table wizard as described in this topic. If you copy an existing table from another IT data mart or job and move it into a different job, or if you attempt to reattach a staged table to its original staging transformation after it was removed, then you might lose valuable components of the table, such as computed columns.

To create the metadata for a staged table, perform the following steps:

1. Open a staging job that includes a staging transformation. For more information about creating a staging transformation, see “Add a Staging Transformation to a Staging Job” on page 109.
2. Right-click the staging transformation in the process flow diagram and select **Add New Staged Table**. The Staged Table wizard opens and prompts you to enter the information that is required to create a staged table.
3. Specify an IT data mart, storage location, and formulas.

Display 7.5 Specify IT Data Mart and Formulas Page of the Staged Table Wizard

- a. Confirm and specify the IT data mart and the location where the staged tables will be stored. The **IT Data Mart** and **Location** fields default to the IT data mart name and location of the current staging job. You must specify a location for the staged table that is within the IT data mart that is specified in the **IT Data Mart** field.

You can use the **Browse** option to select a different location for the staged table or to create subfolders inside the IT data mart for storing the staged tables. However, the Adapter Setup wizard creates and stores staged tables in the same

folder as the staging job so you might choose to accept the default values for consistency.

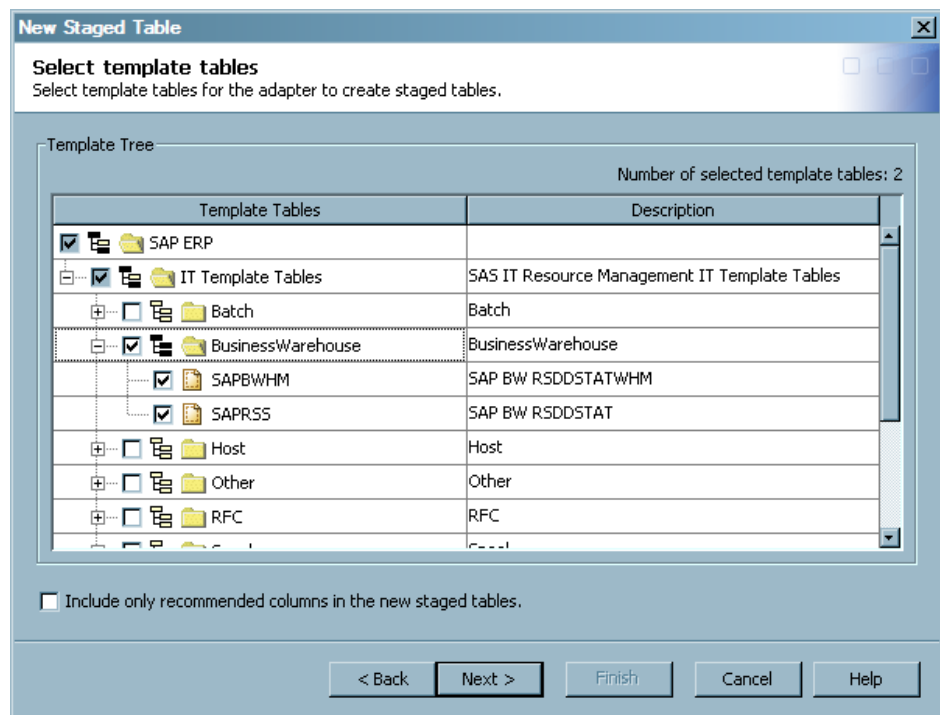
- b. Specify the location for any user-written formulas that you want the Staged Table wizard to use in addition to those supplied by SAS IT Resource Management. The default location directs the Staged Table to the supplied formulas. This field does not allow manual entries. Use **Browse** to navigate to the appropriate folder as needed.

The default location for the formulas folder is the **IT Formulas** folder, which is located here: **Shared Data/SAS IT Resource Management/IT Formulas**.

Note: If you select or create a folder location where the formulas that are required for the staged table are not found in the folder location or its subfolders, the column will not be created. The failure to create the column will be noted in the SAS log file.

- c. Click **Next** to continue.
4. Select template tables. Based on the adapter for the staging transformation, a list of available template tables and their descriptions are displayed within domain categories of table types.

Display 7.6 Select Template Tables Page of the Staged Table Wizard



- a. Use the check boxes to select the template tables that you want to use to generate staged tables. (You can click + to expand a folder or click – to collapse a folder.) If you select a folder, all of the subfolders and template tables within that folder are automatically selected and the check box displays a check (☒) next to the folder and its template tables. If you select an individual template table within a folder, but not all template tables within the given folder, then the check box for the parent folder shows a darkened box with a check (☒) . This indicates that some, but not all, template tables within the folder are selected.

The **Number of selected template tables** maintains a real-time count of the template tables that you select.

CAUTION:

The more template tables that you select and staged tables that you add to a job, the longer it takes to create the staged tables and later run the job. If you select more than 100 template tables, then a warning message explains that creating a large number of staged tables at once can result in long processing times or processing errors.

- b. Select **Include only recommended columns in the new staged tables** if you want only the recommended columns to be included in the staged tables that you are creating. Recommended columns are those that SAS IT Resource Management has predefined as critical or relevant for reporting and analysis in the performance area of the staged table. If this option is not checked, the staged tables will include all available columns.

Note: You can view the recommended columns in a template table from the **Columns** tab of the Properties dialog box for a given template table. For more information about how to access and view the recommended columns of a template table, see [“Columns in Template Tables” on page 79](#).

- c. Click **Next** to continue.

5. Specify the SAS library parameters for the staged tables and their staging transformation.

Note: This page appears only if you are creating the first staged tables for a staging transformation. If you have already created staged tables for the staging transformation, then the wizard skips this page. Library parameters are set once for a staging transformation and all staged tables for that transformation use the same library parameters. However, if all staged tables are later removed from the staging transformation and new staged tables are subsequently added, then the Specify SAS Library page is displayed again in the wizard. In this case, the staged tables in the subsequent iteration are placed in a different library from the first set of tables.

Display 7.7 Specify SAS Library Page of the Staged Table Wizard

New Staged Table

Specify SAS library
Create a new SAS library

SAS library specification

Name: SAPR3 Staging 27415

Description: Contains all staged tables connected with "SAP ERP" transformation in "SAPJob1" job.

Server: SASITRM

Libref: STG27415

Path Specification: C:\SAPJob1\stg27415

< Back Next > Finish Cancel Help

- Specify the **Name** and **Description** for the SAS library. The SAS library name must be unique within its folder in the application server. The Staged Table wizard uses the naming convention **<adapter name> + Staging <unique number>** to name SAS libraries. You can accept the default name or enter a new name for the library. However, library names must be unique within the application server.
- Confirm that the **Server** field shows the server that is associated with the selected IT data mart where this library is to be stored. This field does not enable modifications.
- Specify the **Libref** that is associated with the staged table library. The Staged Table wizard uses the naming convention **STG + <the same number generated for the library>** to name the libref. You can accept the default name or enter a new libref name.

Note: The following naming rules apply to the libref:

- The libref can consist of no more than eight characters.
 - The libref must begin with a letter (A through Z) or an underscore (_).
 - The remaining characters of the libref must be letters (A through Z), numerals (0 through 9), or an underscore (_). The libref cannot contain spaces.
- Specify the **Path Specification** for the location of the library and thus where the tables for this library are stored. The default value is based on the combination of the default path that is specified for the IT data mart, and the operating system of the server. The characters **stg<unique number>** are also appended to the path, where the unique number is the same as that generated for the preceding library.

For example, if the default path for an IT data mart on Windows is **C:\SAPJob1**, then the default value for a new staging library might be **C:\SAPJob1\stg27416**, as shown. If the default path for an IT data mart on

UNIX is `/u/<user ID>/Sample52`, then the default value for a new staging library might be `/u/<user ID>/Sample52/stg27416`.

You can accept the default path, enter a new path, or click **Browse** to select a path.

Note: **Browse** is disabled when the selected application server is running on a machine using the z/OS operating system. When entering a z/OS path manually, you can use either a prefix for a traditional z/OS filesystem path (**MY.DATAMART**) or a root directory in the UFS hierarchical file system (`/u/myname/datamart`).

- e. Click **Next** to continue. The completed Staged Table wizard will create one SAS library that stores all of the tables and catalogs from the staging transformation. This includes staged tables, control tables for data duplication, and the compiled macro code for the transformation.
6. Specify the z/OS attributes that allocate file space in the operating environment if the application server is running on z/OS. The Staged Table wizard skips this page for operating environments that are not z/OS. For more information about these parameters, see the Help for SAS IT Resource Management.
7. Review the details of the staged tables and storage parameters that you entered.

Display 7.8 Summary Page of the Staged Table Wizard

☐ ☐ ☐ New staged tables will be created with the following details:

```

Adapter name: SAP ERP
IT data mart: MyIT Data Mart
Application server: SAS
Location: /Shared Data/SAS IT Resource Management/IT Data
Marts/MyIT Data Mart/SAP ERP
Formulas location: /Shared Data/SAS IT Resource Management/IT
Formulas
Staged tables: 1)SAPBWMM, 2)SAPRSS
Only recommended columns: No

New library specifications:
Library name: SAPR3 Staging 46866
Library description: Contains all staged tables connected with "SAP
ERP" transformation in "SAP Job" job.
Library ref: STG46866
Library path: C:\SAPJob\stg46866

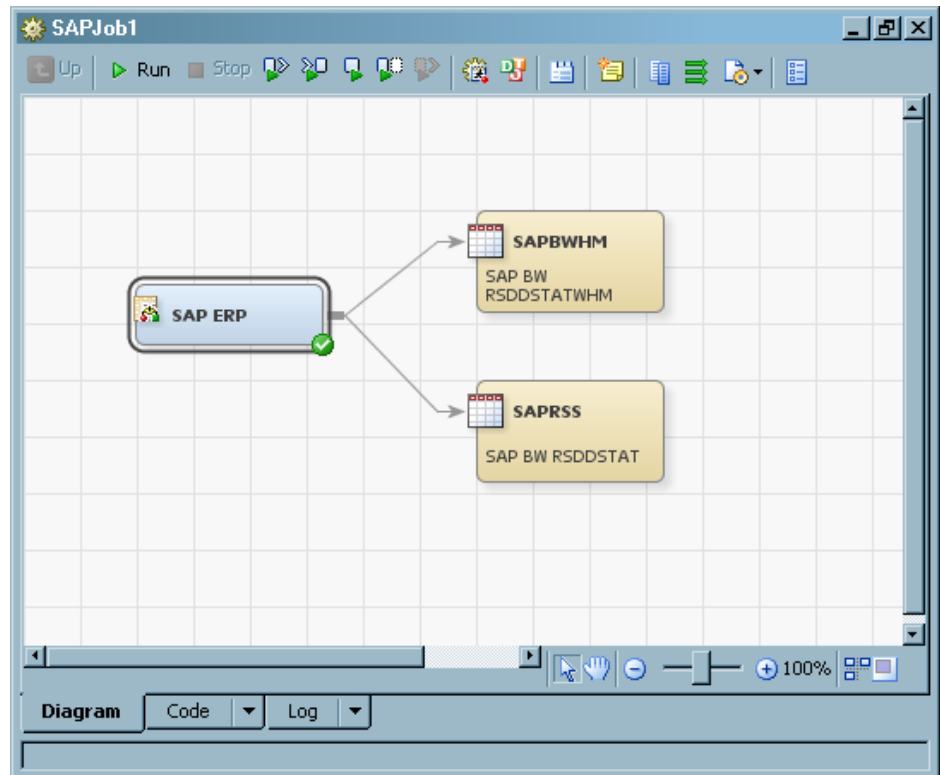
-----
Host: SE
Port: 85
User: ma
  
```

< Back Next > Finish Cancel Help

If you want to change anything that you entered, click **Back** to return to the parameter that you want to change.

Note: You cannot change the name of a staged table because it is automatically assigned the same name as the template table on which it is modeled.

8. Click **Finish** to create the metadata for the new staged table and SAS library (if applicable). The staged tables appear in the process flow diagram for the staging job.

Display 7.9 Process Flow Diagram for Staged Tables in a Staging Job

The staged tables in the process flow diagram show the name and the description of the staged table.

9. Select **File** ⇒ **Save** to save the changes to the job.

Note: The metadata for the staged tables is saved in the staging job. However, the physical staged tables are created only after the staging job is run.

Add an Existing Staged Table to a Job

SAS IT Resource Management does not enable you to create a staged table without associating it to a staging transformation. For best results, always create or add staged tables to a job from the Staged Table wizard. If you copy an existing table from another IT data mart or job, copy or move it into a different job, or attempt to reattach a staged table to its original staging transformation, then you might lose valuable components of the table, such as computed columns. For more information about using the Staged Table wizard to add a staged table to an existing job, see [“Create Staged Tables” on page 114](#).

Erase a Staged Table

Erasing a staged table deletes the physical table, its contents, and its metadata.

CAUTION:

If you erase a staged table that serves as a source table in other jobs, then these associated jobs might not execute because they are missing a source table. Simply erasing a staged table does not affect the metadata of the associated aggregation or information map jobs. However, if you erase a staged table and then re-run the staging job, then the data for the deleted staged table is not created. Thus, the jobs that used that table will fail due to a missing source table.

To erase a staged table, perform the following steps:

1. In a tree view of SAS IT Resource Management, locate the staged table that you want to erase.
2. Right-click the staged table that you want to erase and select **Erase**. The staging job that includes the staged table must be closed and currently not in use before erasing.
3. In the confirmation dialog box, click **Yes** to erase the staged table. The staged table disappears from the tree view and the process flow diagram for the job.

Delete a Staged Table from a Tree View

Deleting a staged table from the tree view in SAS IT Resource Management removes the metadata object. This action does not delete the physical data set or the library that is associated with the staged table.

CAUTION:

If you delete a staged table that serves as a source table in other jobs, then these associated jobs might not execute because they are missing a source table. Simply deleting a staged table does not affect the metadata of the associated aggregation or information map jobs. However, if you delete a staged table and then re-run the staging job, then the data for the deleted staged table is not created. Thus, the jobs that used that deleted table will fail due to a missing source table.

To delete a staged table from a folder in a tree view of SAS IT Resource Management, perform the following steps:

1. In a tree view of SAS IT Resource Management, locate the staged table that you want to delete.
2. Right-click the staged table that you want to delete. The staging job that includes the staged table must be closed and currently not in use before deleting.
3. Select **Delete**.
4. In the confirmation dialog box, click **Yes** to delete the staged table. The staged table disappears from the tree view and the process flow diagram for the job.

Note: If you modify or delete a physical table without using the process described in this topic, the metadata for the table will not be updated.

Delete a Staged Table from a Job

Deleting a staged table from a job removes it from the job so that it is no longer associated with the transformation in the job. This action removes the staged table from the job but does not delete the table's metadata, the physical table, or the library that is associated with the staged table.

CAUTION:

If you delete a staged table that serves as a source table in other jobs, then these associated jobs might not execute because they are missing a source table. Simply deleting a staged table from a job does not affect the metadata of the associated aggregation or information map jobs. However, if you delete a staged table and then re-run the job, then the data for the deleted staged table is not created. Thus, the jobs that used the deleted staged table (or any existing tables that used the deleted staged table) will fail due to a missing source table.

To delete a staged table from a job, perform the following steps:

1. In the IT Data Marts tree view of SAS IT Resource Management, locate the job that includes the staged table that you want to delete.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
3. Right-click the staged table object in the process flow diagram and select **Delete**. The staged table is no longer visible in the process flow diagram of the job.

Purge the Content of a Staged Table

To purge the contents of a staged table while keeping the physical table and its metadata intact, perform the following steps:

1. In the IT Data Marts tree view of SAS IT Resource Management, locate the job that includes the staged table that you want to purge.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
3. Right-click the staged table object in the process flow diagram and select **Purge**.
4. In the confirmation dialog box, click **Yes** to purge the staged table.

Modify a Staged Table

To modify a staged table, navigate to and open the job that contains the staged table that you want to modify.

CAUTION:

For best results, always modify staged tables from the process flow diagram. If you access a staged table's Properties dialog box by opening a staged table from a tree view in SAS IT Resource Management, then you might not see all modification options that are available for a staged table.

Perform the following steps:

1. In the process flow diagram of the job, right-click the staged table that you want to modify.
2. Select **Properties** to open the Properties dialog box.
3. On the **General** tab, you can change the name and description of the table.
4. On the **Columns** tab, you can perform the following tasks:

Add a New Data Column to the Staged Table

- a. Select **New** ⇒ **Data Column**. A new line is added to the grid.
- b. Click the corresponding cells in the new grid line to enter the name (must be unique), description, type, length, informat, and format for the new data column.


Note: The **Expression** cell is not available because the new column is a data column and not a computed column.

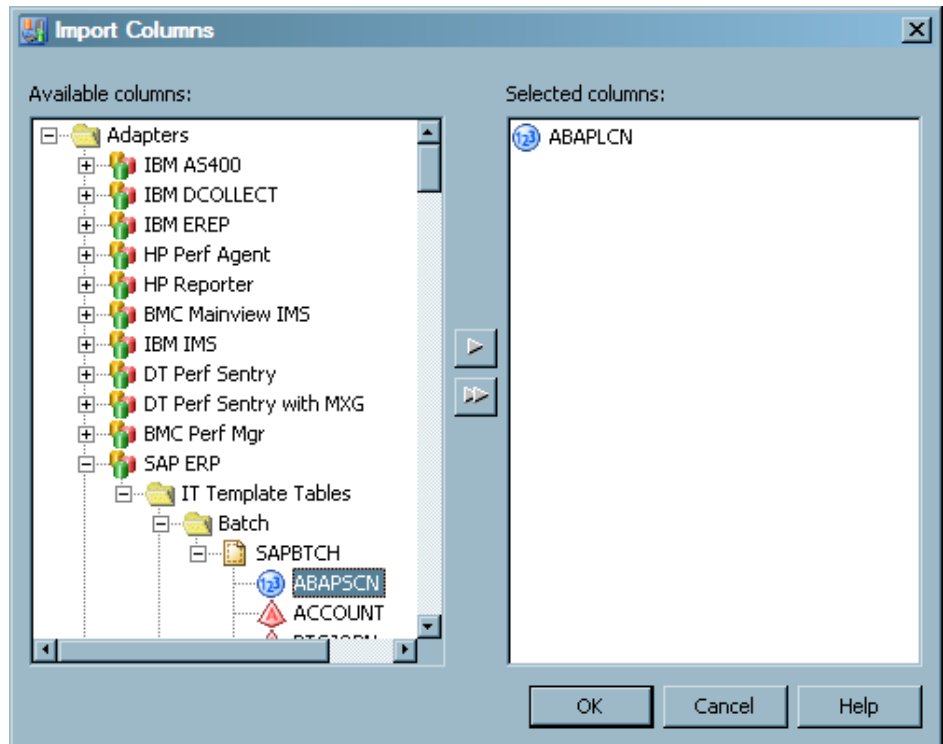
Add a New Computed Column to the Staged Table

- a. Select **New** ⇒ **Computed Column**. A new line is added to the grid.

- b. In the new grid line, enter the name (must be unique), expression, description, type, length, informat, and format for the new data column.

Import a Column from Another Table into the Staged Table

- a. Click  to open the Import Columns dialog box.



- b. In the **Available columns** list, navigate to the table that contains the column that you want to import.

Note: "User" is a reserved word in SQL and should not be the name of any imported column.

- c. Highlight the column that you want to import.
- d. Click the right arrow to transfer the column into the **Selected columns** list.
- e. Click **OK** when you are satisfied with your list of selected columns to import. The new columns appear in the grid.

Note: After you import the column, you must run the staging job before you can view the data in the new column. When the staging job runs the data is propagated to the physical table and is available for viewing.

Modify the Properties of a Column in the Staged Table

- a. Locate the column that you want to modify and double-click the cell that you want to change.
- b. Enter the new name (must be unique), expression, description, type, length, informat, or format value for the column. The **Expression** cell is not available for data columns. The **External Name** cell is not applicable for all columns.

Note: You can also use the icons in the top row of the **Columns** tab to manage, move, and modify additional properties of each column of the staged table.

For more information about using these features, click **Help** on the **Columns** tab.

Delete a Column in the Staged Table

- a. Locate the column that you want to delete and right-click the column name.
- b. Select **Delete**. The column is immediately removed from the grid and no confirmation message is displayed.

Note: If you delete an instanced variable from a staged table, then the staging code might not run successfully later.

5. On the other properties tabs (**Indexes**, **Keys**, **Physical Storage**, **Options**, **Notes**, **Extended Attributes**, and **ITRM Options**), you can view additional properties for the staged table. For more information about the parameters on these tabs, click **Help** on each tab.

Note: The values displayed on the **Keys** tab are ignored by SAS IT Resource Management.

6. When you are finished, click **OK** to save your changes.

Columns in Staged Tables

The **Columns** tab is available from the Properties dialog box of a staged table. It enables you to view the specific column metadata for the table. You can use this information to view the types of data that the staged table includes and understand how computed columns are formulated.

The following list describes the information that you can view for each column in a staged table. Each row on the **Columns** tab shows the metadata for a single column of the staged table.

Note: To access the Properties dialog box and **Columns** tab of a staged table, you can right-click a staged table in the tree view or in the process flow diagram. However, the properties that are available differ with each method. For best results, access the properties from the staged table object in the process flow diagram.

#

specifies the column's ordinal position in the grid on the **Columns** tab.

Name

specifies the name of the column in the staged table. The name of a computed column can exceed eight characters.

Expression

specifies the combination of functions and mathematical operations that are used to derive a value for the column. The field is blank unless the column is a computed column. For more information about computed columns in staged tables, see [“Computed Columns in Staged Tables” on page 126](#).

This field is available only when opening the Properties dialog box from the process flow diagram.



Description

provides a description of the data that populates the column. The description provides more information about the type of data that the column includes.

The column descriptions in staged tables are used as labels in information maps and aggregations.

Type

specifies the data type of the column. Valid types are **Character** and **Numeric**.

You can also identify column types by looking at the Name variable. A round  icon to the left of the name indicates that the column is numeric. A pyramid  icon to the left of the name indicates that the column contains character data.

Length

specifies the maximum character length of the data in the column.

Informat

specifies the SAS informat that is used (if needed) to read values for the selected column. Informats are required for reading fixed-width data.

Format

specifies the SAS format that is used to write values for the selected column. If the staged table column was instantiated from a template table, the format is in NLS format when applicable. NLS formats are indicated by NL as the first two characters in the format column. For more information about NLS formatting, see [“National Language Support \(NLS\)” on page 12](#).

External Name

specifies the name of the corresponding field in the raw data that is used to populate this column. External names are the means by which the staging code maps columns in a staged table to their corresponding variables in the raw data. Not all columns in a staged or template table have a predetermined external name. A value is optional in this field for these columns. However, if the column is associated with an adapter that uses external names, do not change an external name that is provided. For more information about external names, see [“Understanding External Names of Staged Tables and Columns” on page 127](#).

Note: This field is available only when opening the Properties dialog box from the process flow diagram.

Is Nullable

indicates whether a column can contain null or missing values.

This field is available only when opening the Properties dialog box from the tree view.

Summary Role

indicates the intended role of the column in summary data.

This field is available only when opening the Properties dialog box from the tree view.

Sort Order

defines how the values in the column are sorted, such as ascending or descending.

This field is available only when opening the Properties dialog box from the tree view.

Note: The **Data Column**, **Computed Column**, or **Import Columns** options might be disabled if you used the up or down arrows to change the sort order of the columns. To re-enable these options, click the # icon that appears in the grid of columns.

Computed Columns in Staged Tables

Computed columns in a table contain data values that are derived from other columns in the table. Staged tables contain computed columns for data that is not directly provided by the raw data source but the data might be useful for reporting or promoting consistency across various data sources. For example, a data source normally has a datetime stamp and from that datetime stamp, SAS IT Resource Management might provide a computed column that calculates weeks or months.

Note: A computed column should not be used as a source of calculations for other computed columns. Using a computed column as a source for a computation can be problematic because the order of computing the values for the computed column is not guaranteed.

To find a computed column in a staged table and view the formula that it uses, right-click the staged table in a process flow diagram and select **Properties**. Then, open the **Columns** tab of the staged table's Properties dialog box. This method is the only way to view the **Expression** column.

Note: If you open a table from a folder view or from the IT Data Marts view, and select the **Columns** tab, the columns of the table are displayed. However, the expressions that are associated with the columns are not displayed.

All computed columns contain a formula or expression in the **Expression** column of the grid. You can double-click an expression to modify it for the column. To open the Expression Builder, click the ellipsis button in the cell. For information about how to work with the Expression Builder, click **Help** in the Expression Builder window.

Note: Aggregation tables also use computed columns. You can use the computed columns that are provided for these tables and you can create your own. For more information about creating and using computed columns, see [“Specify Computed Columns” on page 153](#).

Staging VMware Data Using Local Time

About Staging VMware Data Using Local Time

All datetime stamps are stored in Coordinated Universal Time (UTC) or Greenwich Mean Time (GMT) in the VMware vCenter database. However, you can convert the datetime-related values in your staged tables (and thus in your IT data mart) to use local time instead of UTC or GMT.

Note: The following procedures apply only to the data processed by the VMware vCenter adapter.

To change the datetime-related values in your IT data mart from UTC to local time, perform the following steps:

1. Change the formula expression for the IT Formula VMwareDatetime.
2. Change the expressions for other datetime-related computed columns in your staged tables. To do so, change the computed column expression from **<datetime-**

```
related variableName> + 0 to <datetime-related variableName>
+ GMT offset.
```

Change the Formula Expression for the IT Formula VMwareDatetime

To change the values for the DATETIME column in your VMware staged tables from UTC to local time, modify the expression for the IT Formula, VMwareDatetime, by adding the GMT offset to the formula expression.

By changing the expression for the VMwareDatetime IT Formula, the datetime value adjustments will be made during staging for the DATETIME column in all VMware staged tables. For example, to adjust the DATETIME variable from UTC to Eastern Daylight Time in your staged tables, change the expression for the VMwareDatetime IT Formula from `Datetime + 0;` to `Datetime - 14400;`.

Change the Expressions for Other Datetime-related Computed Columns in Your Staged Tables

To change the other datetime-related values in your staged tables from UTC to local time, change the expressions for those datetime-related computed columns in your staged tables. To do so, change the computed column expression from `<datetime-related variableName> + 0` to `<datetime-related variableName> + GMT offset`. For example, to adjust the CREATETIME variable in the Events staged table from UTC to Eastern Daylight Time during staging, change the computed column expression from `CREATETIME + 0;` to `CREATETIME -14400;`.

Note: If you modify any one of the datetime-related columns to be adjusted with a GMT offset, then all of the datetime-related columns in your staged tables should be adjusted with the same GMT offset.

The following list shows the datetime-related values that can be changed for each type of VMware staged table:

- All VMware staged tables - DATETIME variable
- EVENTS staged table - CREATETIME
- TASKS staged table - QUEUE TIME, STARTTIME, COMPLETETIME

Understanding External Names of Staged Tables and Columns

About External Names of Staged Tables and Columns

Staged tables and the columns in a staged table can each have an external name. These external names are the means by which the staging code maps the SAS IT Resource Management staged tables and columns to the tables and variables in the raw data. For example, a staged table that is named NTPHDSK in SAS IT Resource Management has an external name of PhysicalDisk. The staging code uses this information to determine that the PhysicalDisk data in the raw data file belongs in the NTPHDSK staged table. The same is true for columns and the external names for their corresponding variables.

Note: The SAR adapter is the only supported adapter that does not require external names.

You can view the external names that are associated with a staged table and its columns from the Properties dialog box of the staged table.

CAUTION:

For best results, do not modify the external names of a staged table or column that is associated with an adapter that uses external names. Changing the external names that are associated with adapters that use external names can cause errors when the raw data is staged. The staging code for these adapters relies on external names to associate the raw data file with the appropriate staged tables and columns. However, you can modify the external names of staged tables and columns that are associated with the SAR adapter, because this adapter does not use external names. The SAR adapter ignores the values for external names, so there is no danger of corrupting the staging metadata.

Change the External Name of a Table or a Variable

CAUTION:

Do not change the external name that is prepopulated for a staged table or variable unless instructed to do so by a SAS representative. You should change a table external name or variable external name only if the data source has changed the table or variable name that is output in the raw data, and if the SAS IT Resource Management table or variable external name has not been updated to reflect that change.

For most adapters, SAS IT Resource Management uses the external names in order to map the tables and variables in the raw data to the SAS IT Resource Management table and column names. For example, previous versions of the Demand Technologies Performance Sentry collector wrote an object to the log file as **SQLServer-Locks**. However, in later versions, this object name changed to **SQLServer:Locks**. In order to process data for the new object, the external name value for the SAS IT Resource Management table NTSLOCKs should be changed from **SQLServer-Locks** to **SQLServer:Locks**. In the preceding example, the mapping could not take place and the updated table or variable could not be processed.

If you expect to process log files that will contain only the new table or variable external names, then you can resolve the situation by opening the table properties and replacing the existing table or variable external name with the new table or variable external name. (This is a rare situation and you should change external names only when this particular condition occurs.)

View the External Names of Tables

To view the external name for a staged table, perform the following steps:

1. Right-click a staged table in the tree view or in a process flow diagram.
2. Click **Properties** to open the Properties dialog box.
3. Select the **Extended Attributes** tab. The **ExternalName** field shows the external name that is assigned to the staged table.

Note: Do not change the external name for a staged table unless instructed to do so by a SAS representative.

View the External Names of Columns

To view the external name for a column in a staged table, perform the following steps:

1. Right-click a staged table in the process flow diagram.
2. Click **Properties** to open the Properties dialog box.
3. Select the **Columns** tab. The **External Name** field provides the external names of the variables that are associated with the columns in a staged table. Not all columns have an external name that is predetermined. You can add a value in the **External Name** field for these columns.

Note: Do not change an external name that is prepopulated for a column unless instructed to do so by a SAS representative.

External Names for MXG Adapters

For MXG adapters, the external table name is the name of the SAS data set that MXG creates. The external column names are the variable names that MXG uses within those data sets. There is no benefit to changing the external names of the tables or columns for MXG adapters.

CAUTION:

Do not change the external names of the columns or the tables for MXG adapters. If you do change these external names, your data could be corrupted and errors could occur in the processing of that data.

For more information about working with MXG data and adapters, see [“Data Sources Supported by SAS IT Resource Management Adapters” on page 290](#).

Insufficient Disk Space

If there is not enough disk space to run a staging job, then the job fails and a message indicates that there was a failure writing to the disk. If this occurs, then point the failing SAS library to an alternate path that has more space available or delete unwanted files on that disk. Then, rerun the staging job.

Chapter 8

Aggregating the Data

About Aggregation Tables	132
What Is an Aggregation Table?	132
How Is the Metadata for an Aggregation Table Created?	132
How Is the Physical Aggregation Table Generated?	134
What Are the Properties of an Aggregation Table?	134
Working with Aggregation Tables	140
Set up the Aggregation Transformation in the Process Flow Diagram	140
Creating Aggregation Tables with the Summarized Aggregation Table Wizard	141
How to Enter General Information about the Aggregation Table	141
How to Specify Library Information	141
How to Select a Filter	142
Specify Aging	142
How to Specify Class and ID Variables	144
How to Specify Column Details	145
How to Specify Analysis Variables	145
Specify Statistics	146
Specify Percent Change	149
Specify Ranking	150
Specify Join Columns	151
Specify Computed Columns	153
Completing the Specification of the Summarized Aggregation Table	157
Creating an Aggregation Table with the Simple Aggregation Table Wizard	157
Creating Aggregation Tables with the Aggregations through Templates Wizard	159
Modifying Aggregation Tables	160
Delete an Aggregation Table from an Aggregation Transformation	160
Delete an Aggregation Table from the IT Data Mart	161
Delete an Aggregation Table with Contents	161
Erase an Aggregation Table	161
Modify a Simple or a Summarized Aggregation Table	162
Purging the Content of an Aggregation Table	164
Renaming an Aggregation Table	164
Working with Aggregation Transformation Filters	165
About Aggregation Transformation Filters	165
Access the ITRM Options Tab on the Properties Dialog Box of an Aggregation Transformation	165
Define a Filter for an Aggregation Transformation	166
Delete a Filter from an Aggregation Transformation	166

Delete a Filter from an Aggregation Table	167
Modify a Filter of an Aggregation Transformation	167
About the Expression Builder	168
Working with Aggregation Transformations	170
Properties of an Aggregation Transformation	170
Delete an Aggregation Transformation from a Job	173
Rename an Aggregation Transformation	173
Indexing an Aggregation Table	173
Introduction to Indexes	173
How to Access an Index	174
Create an Index	175
Delete a Column from an Index	175
Delete an Index	176
Rename an Index	176

About Aggregation Tables

What Is an Aggregation Table?

Aggregation tables contain data that is collected and accumulated according to rules that are specified by an Aggregation transformation. These rules are based on columns from the source table and are used to classify, organize, and calculate statistics about IT data. The target table that is produced by an Aggregation transformation is used for further data management, analysis, or report processes.

With SAS IT Resource Management, you can generate two types of aggregation tables: summarized aggregation tables and simple aggregation tables.

- Summarized aggregation tables:

Summarized aggregation tables consist of data from the source table that is summarized in accordance with the rules that are specified in the Aggregation transformation. Summarized aggregation tables also include statistics that are calculated and stored for the categories of data that you specify.

- Simple aggregation tables:

Simple aggregation tables consist of data from the source table that is simply appended to the target table. Like summarized aggregations, simple aggregations support join columns, computed columns, and aging. Simple aggregations are analogous to the detail tables of the performance data warehouse (PDB) in SAS IT Resource Management 2.7.

How Is the Metadata for an Aggregation Table Created?

Methods for Creating Metadata for an Aggregation Table

The metadata for an aggregation table can be created in two ways:

- automatically, by means of the Adapter Setup wizard
- directly, by invoking an aggregation wizard from an Aggregation transformation

Creating Metadata with the Adapter Setup Wizard

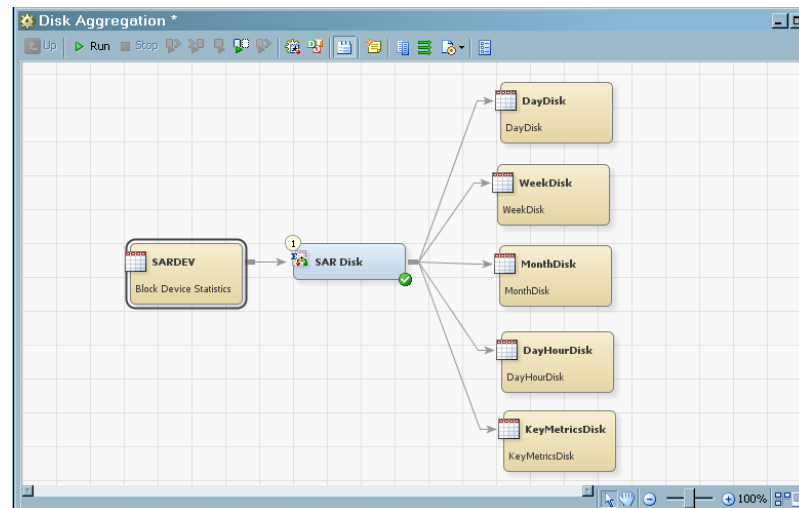
Using the Adapter Setup wizard, you can specify the adapter, domain category, and time periods for which you want to generate reports. The Adapter Setup wizard creates the metadata for the appropriate Aggregation transformations, according to your specifications.

To invoke the wizard, from the menu bar, select **New** ⇒ **IT Resource Management** ⇒ **Adapter Setup**.

Note: The Adapter Setup wizard prompts for other input, such as the IT data mart where the job is to be stored, the location of the formulas that should be used when creating computed columns, and information about where the reports are to be stored. For more information about the Adapter Setup wizard and the other objects that it creates, see [“About the Adapter Setup Wizard” on page 177](#).

The following display shows the Disk job that was generated by the Adapter Setup wizard for the SAR adapter. It contains a SAR Disk Aggregation transformation that specifies the creation of five aggregation tables.

Display 8.1 SAR Disk Job That Generates Multiple Aggregation Tables



Note: The Adapter Setup wizard can generate multiple aggregation tables from a single Aggregation transformation.

Creating Metadata with an Aggregation Wizard

To create an aggregation table with the aggregation wizard from an Aggregation transformation, you must first perform the preliminary steps. For information about how to perform these steps, see [“Set up the Aggregation Transformation in the Process Flow Diagram” on page 140](#).

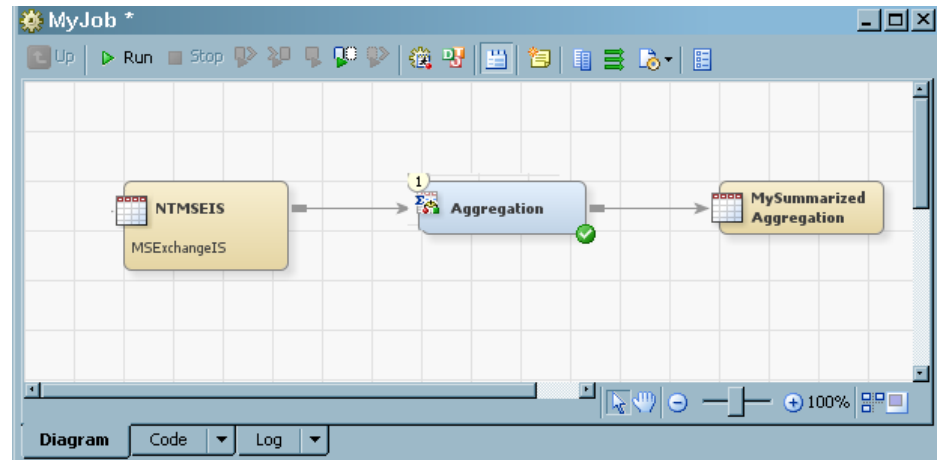
To invoke the wizard, right-click the Aggregation transformation and select **Add Aggregation Table** from the drop-down list. You can then choose to create a **Summarized Aggregation Table**, a **Simple Aggregation Table**, or **Aggregations through Templates**.

You must respond to the prompts of the selected wizard in order to specify how the aggregation table is to be generated. When the wizard is successfully completed, the metadata for the aggregation table is generated and stored in the folder of the IT data mart that you specified. (If you do not select an IT data mart, the wizard defaults to the IT data mart in which the job resides.) The **Diagram** tab of the Job Editor window is

updated to display the job that, when run, generates the aggregation table that you specified.

The following display shows a job called MyJob that was generated by directly invoking the Summarized Aggregation wizard. It contains an Aggregation transformation that is specified to create a single aggregation table.

Display 8.2 Table Generated by the Summarized Aggregation Wizard



Note: An Aggregation wizard that is directly invoked can generate only one aggregation table for each invocation of the wizard. However, you can invoke the wizard multiple times on the same Aggregation transformation, each time generating another aggregation table

How Is the Physical Aggregation Table Generated?

Aggregation transformations are contained in jobs that generate physical aggregation tables when the jobs are successfully run. For information about how to run a job immediately, see [“Running Jobs” on page 273](#). For information about how to deploy a job for batch scheduling, see [“Run a Job Immediately” on page 273](#).

What Are the Properties of an Aggregation Table?

To open the Properties dialog box of an aggregation table, right-click the aggregation table in the process flow diagram. From the drop-down menu that is displayed, select **Properties**.

The following display shows the Properties dialog box for the DayDisk aggregation table, which was displayed by right-clicking the DayDisk aggregation table image on the

process flow diagram. (The metadata for this table is one of many tables that are generated by the Disk Aggregation transformation of the SAR adapter.)

Display 8.3 Properties for the DayDisk Aggregation Table

The screenshot shows the 'DayDisk Properties' dialog box with the 'General' tab selected. The 'Name' field contains 'DayDisk'. The 'Description' field contains 'DayDisk'. The 'Location' field shows a file path: '/Shared Data/SAS IT Resource Management/IT Data Marts/My IT Data Mart/SAR 1/Domain Categories/Disk'. There is a 'Browse...' button next to the location field. Below the location field is a table for 'Responsibilities' with two columns: 'Name' and 'Role'. The table is currently empty. At the bottom right of the dialog are buttons for 'New', 'Delete', 'OK', 'Cancel', and 'Help'.

Note: Unlike other SAS Data Integration Studio transformations, there is no **Mapping** tab. The aggregation wizard manages this function.

The Properties dialog box consists of the following tabs:

- The **General** tab contains the identifying information about the table. The fields on this tab can be modified.

The following display shows the **General** tab of an aggregation table.

This is a duplicate of the screenshot above, showing the 'DayDisk Properties' dialog box, General tab. It displays the Name, Description, and Location fields, an empty Responsibilities table, and the bottom action buttons (New, Delete, OK, Cancel, Help).

For items in the preceding display, the following definitions apply.

Name

identifies the name of the aggregation table.

Description

describes the aggregation table.

Location

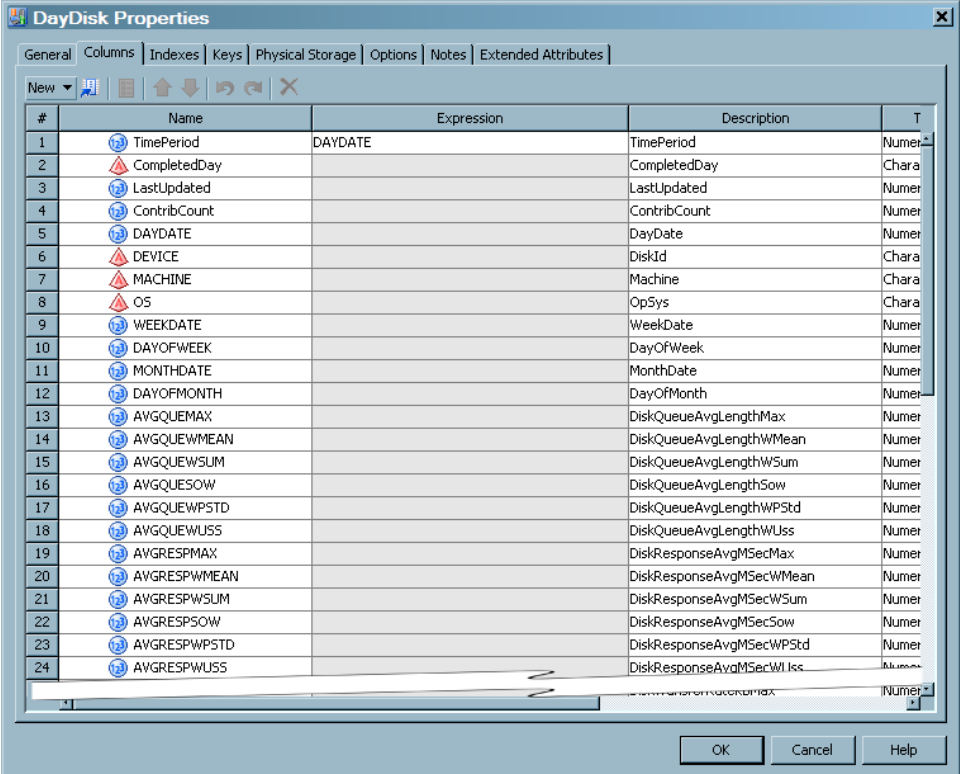
specifies the location of the aggregation table in the metadata folder hierarchy.

Click **Browse** to select a path to the table from the Select a Location dialog box that displays the metadata directory structure.

- The **Columns** tab contains a grid that provides the following information about each column of the table. The Aggregation wizards supply the following columns:
 - **ContribCount** is a numeric column in summarized aggregation tables. The value of ContribCount is the number of input observations (rows of data) that were used in computing this summarized row of data
 - **LastUpdated** is the system datetime of when this row of data was last updated. This is in simple and summarized aggregation tables.
 - **CompleteDay** is a single character, set to 'Y', 'N', or blank. If the Completed Day processing is not enabled, it is always set to blank. If the Completed Day processing is enabled, this column is set to 'Y' if the value of DATETIME indicates that this row of data represents a completed day. Otherwise, this column is set to 'N', which indicates an incomplete day. A day is considered complete if any input row of data has a time portion of the datetime value that is greater than the specified cutoff time.

The following display shows the **Columns** tab of an aggregation table.

Display 8.4 Columns Tab of an Aggregation Table (partial listing)



#	Name	Expression	Description	T
1	TimePeriod	DAYDATE	TimePeriod	Numer
2	CompletedDay		CompletedDay	Chara
3	LastUpdated		LastUpdated	Numer
4	ContribCount		ContribCount	Numer
5	DAYDATE		DayDate	Numer
6	DEVICE		DiskId	Chara
7	MACHINE		Machine	Chara
8	OS		OpSys	Chara
9	WEEKDATE		WeekDate	Numer
10	DAYOFWEEK		DayOfWeek	Numer
11	MONTHDATE		MonthDate	Numer
12	DAYOFMONTH		DayOfMonth	Numer
13	AVGQUEUMAX		DiskQueueAvgLengthMax	Numer
14	AVGQUEWMEAN		DiskQueueAvgLengthWMean	Numer
15	AVGQUEWSUM		DiskQueueAvgLengthWSum	Numer
16	AVGQUESOW		DiskQueueAvgLengthSow	Numer
17	AVGQUEWPSTD		DiskQueueAvgLengthWPStd	Numer
18	AVGQUEWUSS		DiskQueueAvgLengthWUss	Numer
19	AVGRESPPMAX		DiskResponseAvgMSecMax	Numer
20	AVGRESPPWMEAN		DiskResponseAvgMSecWMean	Numer
21	AVGRESPPWSUM		DiskResponseAvgMSecWSum	Numer
22	AVGRESPPSOW		DiskResponseAvgMSecSow	Numer
23	AVGRESPPWSTD		DiskResponseAvgMSecWPStd	Numer
24	AVGRESPPWUSS		DiskResponseAvgMSecWUss	Numer


For items in the preceding display, the following definitions apply.

#

specifies the column's ordinal position in the grid on the **Columns** tab.

Name

specifies the name of the column.

If the Key icon  appears next to the name, then that column is used as a primary, unique, or foreign key.



Expression

specifies the combination of functions and mathematical operations that are used to derive a value for the column. This field is blank unless the column is a computed column.

Description

describes the column.

Type

specifies the data type of the column. Valid types are character and numeric. You can also identify column types by looking at the Name variable. A round  icon to the left of the name indicates that the column is numeric. A pyramid  icon to the left of the name indicates that the column contains character data.

Length

specifies the length of the column. The length of numeric columns is from 2 through 8. The length of character columns is from 1 through 32,767.

Informat

specifies the SAS informat that is used (if needed) to read values for the selected column. Informats are not used by the Aggregation transformation.

Format

specifies the SAS format that is used to write or display values for the selected column.

Summary Role

specifies the type of the column. It is not used by SAS IT Resource Management.

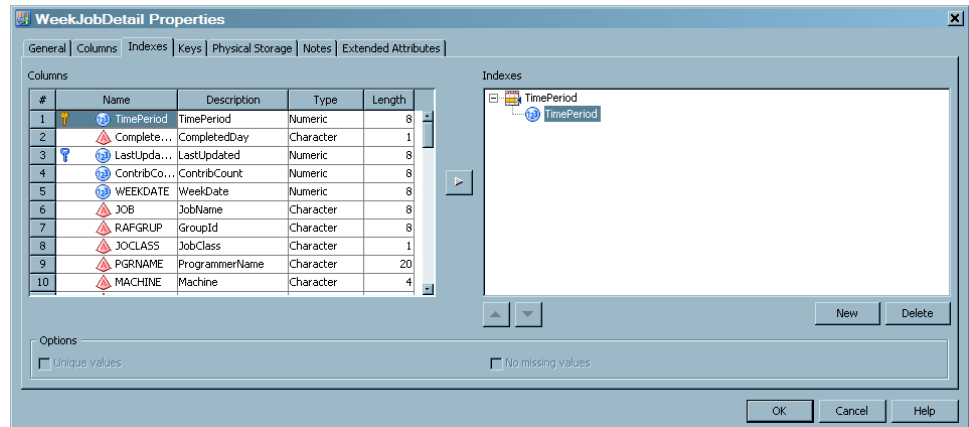
CAUTION:

The New button at the top of the grid should not be used. The preferred method of adding, deleting, or modifying a column is to do so from the appropriate aggregation wizard. Using the wizard ensures that the metadata is properly updated.

- The **Indexes** tab contains a list of the columns (and their descriptions, types, and lengths) that can be selected to use as indexes. Indexes can improve queries that use the aggregation table.

The following display shows the **Indexes** tab of an aggregation table.

Display 8.5 Indexes Tab of an Aggregation Table

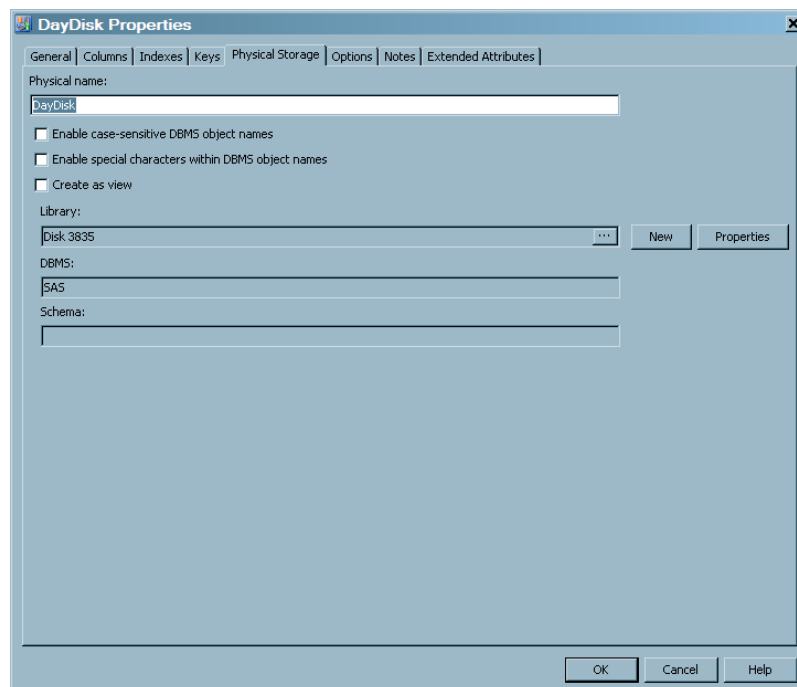


For more information about indexes, see “[Indexing an Aggregation Table](#)” on page 173.

- The **Keys** tab is not used by SAS IT Resource Management, and its values are ignored.
- The **Physical Storage** tab contains information about where the table is stored.

The following display shows the **Physical Storage** tab of an aggregation table.

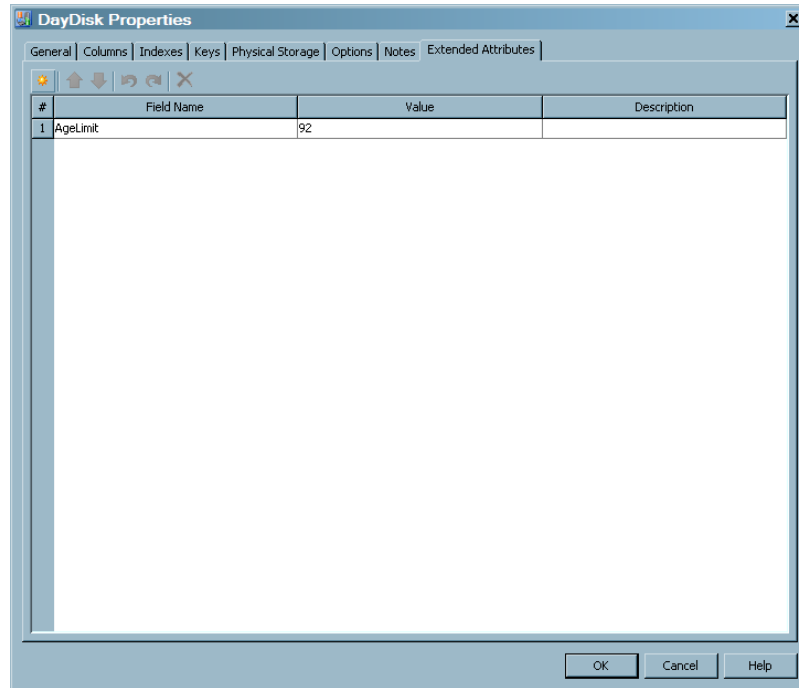
Display 8.6 Physical Storage Tab of an Aggregation Table



- The **Notes** tab contains areas where you can annotate the table.
- The **Extended Attributes** tab contains a custom property that is not part of the metadata for the table.

The following display shows the **Extended Attributes** tab of an aggregation table.

Display 8.7 *Extended Attributes Tab of an Aggregation Table*



For items in the preceding display, the following definitions apply.



adds a row to the attribute table. A number and a default name are provided, but you must enter the value and description.



deletes a selected attribute from the list.



If present, enables you to reorder rows by moving a selected row up in the table.



if present, enables you to reorder rows by moving a selected row down in the table.

#

specifies the number of the extended attribute.

Field Name

specifies the attribute name to be defined.

Value

specifies the value for the attribute that is specified in the **Field Name** column.

Description

describes the attribute and value.

CAUTION:

Do not delete or modify the AgeLimit field from the Extended Attributes tab. In addition, do not delete the Filter from the Extended Attributes tab. Doing so can cause processing errors.

Working with Aggregation Tables

Set up the Aggregation Transformation in the Process Flow Diagram

To create an aggregation table with an aggregation wizard, you must first perform the following preliminary steps:

- Open a job on the **Diagram** tab of the Job Editor window.
- Drag the Aggregation transformation onto the process flow diagram.
- Connect a source table to the transformation.

To perform the preceding preliminary steps, do the following:

1. Locate the job that you want to contain this Aggregation transformation. From the IT Data Marts tree, navigate to the IT data mart and folder that contains the job that you want to use. Double-click the job to open it on the **Diagram** tab of the Job Editor window.

If you want to create a new job, see [“Jobs That Process IT Data” on page 263](#).

2. Navigate to the Transformations tree and expand **ITRM Transformations**.
3. Drag and drop the Aggregation transformation onto the process flow diagram.

Note: A job can contain multiple Aggregation transformations.

4. Drag and drop a source table onto the process flow diagram and using the cursor, connect the table to the left side of the aggregation image.

Note: The source table is usually a staged table, but it can be any SAS table.

5. Right-click the aggregation image and, from the drop-down menu that is displayed, select **Add Aggregation Table**.
6. Select the type of aggregation table that you want to create. (If you click **Cancel** on any page of an aggregation wizard, a confirmation dialog box is displayed. If you do not want to cancel, click **No**.)
 - Select **Summarized Aggregation Table** to invoke the Summarized Aggregation wizard. For information about using this wizard, see [“Creating Aggregation Tables with the Summarized Aggregation Table Wizard” on page 141](#).
 - Select **Simple Aggregation Table** to invoke the Simple Aggregation wizard. For information about using this wizard, see [“Creating an Aggregation Table with the Simple Aggregation Table Wizard” on page 157](#).
 - Select **Aggregations through Template** to invoke the Aggregations through Template wizard. For information about using this wizard, see [“Creating Aggregation Tables with the Aggregations through Templates Wizard” on page 159](#).

Note: You can create summarized aggregation tables multiple times. The aggregation wizard provides a unique suffix for each table that is recreated. Joined columns that are required for the new table are generated from the columns in the new table.

Creating Aggregation Tables with the Summarized Aggregation Table Wizard

To create an aggregation table (that is, the target table) with the Summarized Aggregation Table wizard, respond to the wizard's prompts.

How to Enter General Information about the Aggregation Table

To specify general information about the aggregation table, perform the following steps:

1. Enter the name of the table in the **Name** field. This field can contain up to 60 characters. It must be unique within the metadata folder. It must also be unique among all other table names within the application server. A value is required in this field.

Note: The name of this aggregation table will be displayed on subsequent pages of the wizard.
2. (Optional) Enter the identifying information in the **Description** field. This field can contain up to 200 characters.
3. (Optional) In the **SAS table name** field, you can accept the default value, or enter the physical name of the aggregation table.
4. (Optional) In the **IT data mart** field, you can accept the default value, or enter the name of the IT data mart. The **IT data mart** field contains the job in which the Aggregation transformation is located.
5. (Optional) In the **Location** field, you can accept the default value, or enter the metadata path to the metadata folder where the job that contains the Aggregation transformation is stored.

How to Specify Library Information

If this aggregation table is the first target table that is specified for this Aggregation transformation, an additional page of the wizard opens. This page prompts you for the physical target location and for information about the SAS library where the data will be stored. When this page opens, the following fields are already specified with default values. If you want to modify these fields, perform the following steps:

1. In the **Name** field, enter the name of the library that contains all aggregation tables that are associated with the Aggregation transformation that you are working with. This field can contain up to 60 characters. A value is required in this field.
2. In the **Description** field, enter the information that describes the library. It can contain no more than 200 characters. It is an optional field.
3. In the **Server** field, enter the application server that you specified for the IT data mart that you are working with. The default application server is the SASITRM server. A value is required in this field.
4. In the **Libref** field, enter the name that is temporarily associated with the library. The libref or an aggregation library is a concatenation of AGG and a random number (for example, *AGG27040*). A value is required in this field. The libref name must be unique within the application server.

Note: The following naming rules apply to the libref:

- The libref can consist of no more than eight characters.
 - The libref must begin with a letter (A through Z) or an underscore (_).
 - The remaining characters of the libref must be letters (A through Z), numerals (0 through 9), or an underscore (_). The libref cannot contain spaces.
5. In the **Path specification** field, enter the location where the tables for this library are to be stored. You must specify a location for the tables that is within the IT data mart that is specified in the **IT Data Mart** field.

Enter this field manually or click **Browse** to select a location from the Select a Directory dialog box that displays the directory structure on the server's file system. A value is required in this field.

How to Select a Filter

Filters can limit the amount of data that is input to an aggregation table. To select a filter and a column for data-duplication checking, perform the following steps:

1. From the drop-down list in the **Select filter** field, select the filter that you want to use. The selected filter is displayed in the **Select filter** field and its corresponding expression is displayed in the **Expression** field.
2. The **Expression** field displays the formula that defines the filter. It cannot be changed from this location. For information about how to modify the expression field of a filter, see [“Modify a Filter of an Aggregation Transformation” on page 167](#).
3. Check the **Enable data duplication checking** box to indicate that you want data-duplication checking to be performed as the data from the source table is read. If checked, this option removes duplicate data items from the resulting target table.

Note: This option is enabled only if LSTPDATE is found in the source table. (The default value for this option is deselected.)

If you want to create a new filter to use for your aggregation table, perform the following steps:

1. Click the **Add New** button to open the Add Filter dialog box.
2. In the **Name** field, enter the name of the new filter. The name field can contain no more than 60 characters.
3. In the **Expression** field, enter the formula.
4. Click **OK**.

Note: For more information about how to work with filters, see [“Working with Aggregation Transformation Filters” on page 165](#).

Specify Aging

Managing Your Data: Aging and Purging Options

SAS IT Resource Management provides several ways of managing the amount of data that is kept in the aggregation tables of your IT data mart. You can apply the aging and purging functions to both simple and summarized aggregations. The selection of the

aging and purging attributes for a simple or a summarized aggregation table can be performed on the Specify aging page of the wizard.

- **Purge Aggregated Data**

Purging the data removes the data from the target table before the new data is loaded.

- **Age Aggregated Data**

You can specify whether and how you want to age your data when you create the Aggregation transformation. Based on these specifications, data will be kept in the aggregation tables or aged out of them.

By entering a value for the **Age limit in days**, you are specifying the age range that data will be present in the table. The minimum number of days is 1; the maximum number of days is 9999.

Note: Columns that are used for aging are DATETIME columns. For these columns, one day of aging extends over 24 hours of data collection, which might extend over one or more days.

You are not specifying that the data is to be aged out after that number of days. Instead, the value that you specify is the number of days of data that you want to keep. Calculate what is aged out by subtracting the value of the **Age limit in days** from the latest date of the new or existing data. The resulting date is the date before which data will age out. For example, in the following figure, Day 9 is the latest date of the data and the **Age limit in days** is seven.

Display 8.8 Example of How Data Is Aged out

Day 1	Day 2	Day 3	Day 4	Day 5	Day 6	Day 7	Day 8	Day 9
Data	Data	Data	No data was loaded	Data	Data	Data	No data was loaded	Data

Data is aged out.

Data is kept.

All data from Day 2 and earlier is aged out. All data from Day 3 through Day 9 is kept. However, because no data was loaded on Day 4 and Day 8, only five days of data is kept, namely the data for Day 3, Day 5, Day 6, Day 7, and Day 9.

Note: If you reduce the aging value, the change takes effect the next time the job is run, even if there is no new data.

How to Specify Purging Criteria

To purge the target table before you load new data, check the corresponding check box. This action removes the data that was processed in previous executions of the job.

How to Specify Aging Criteria

You can specify whether and how you want to age your data when you define the aggregation target table. Based on these specifications, data will be kept in the aggregation table or aged out of it. To specify aging criteria for the target table, perform the following steps:

1. To perform aging for the target table, check the corresponding check box. This action enables you to select the column to be used for aging and to specify the format of that column.
 - The **Column** field displays a list of the columns from the source table that are available to use for aging the data. From the drop-down list, select the column

that you want to use. It is displayed in the **Column** field. The format of the selected field is displayed.

Note: If you want to specify percent change for a variable, you must select an aging column. SAS IT Resource Management uses that aging column in order to calculate the period-to-period changes for that variable. The selected aging column is automatically added to the list of class columns for this aggregation table.

- The **Date** or **Datetime** option is automatically selected, based on the contents of the column that you selected for aging. Changing this option is permissible, but it is not recommended.

Note: Specify the option that reflects the type of column that you chose. For DATE values, the aging counts by days. For DATETIME values, the aging counts by the number of seconds in a day.

2. Use the spinner to select the number of days age limit.
3. To specify the completed day, check the corresponding check box. This action enables you to select the column to use to determine whether a day is complete. (This option is enabled if aging has been specified for this aggregation table and the source table has a column called DATETIME.)

This feature is useful only for tables where the aging column represents a DAY or less. For example, this feature is not useful for WEEKDATE aggregations.

Note: If the **Apply completed day** option is deselected on an existing aggregation table that previously selected that option, the rank columns of that table are analyzed. If they depend on Completed Day being selected, a message is displayed that lists the rank columns that would be affected and the deselection of the Completed Day option is not permitted.

4. Use the **Cut off time** field to specify when a day is completed. Use the spinners to specify the hours, minutes, and whether the time is a.m. or p.m.

How to Specify Class and ID Variables

On the Specify class and ID variables page, you can specify the columns from the source table that are to be included in the target aggregation table as class variables and as ID variables. After the class and ID variables have been selected, you can modify the name, label, and the format attributes of the output columns. These changes can be performed on the Specify column details page of the wizard for summarized aggregations.

- A class variable is used to group or classify data. For example, the value of a class variable could be a device address.

At least one variable must be selected as a class variable.

Note: If aging is being performed on this target aggregation table, then the aging column is automatically added to the list of class columns. The aging column cannot be deleted from the list of class variables if it is in use as the aging column. If aging is turned off, the column can be removed from the list of class variables.

- An identification (ID) variable contains an alternate identifier for the data in a class variable. For example, the value of the corresponding ID variable could be the name of a particular device.

ID variables are not required.

To select class columns for the target aggregation, perform the following steps:

1. Highlight the columns in the **Available variables** panel that you want to select as class columns.
2. Click the right arrow to move those columns into the **Selected class variables**.

To select ID columns for the target table, perform the following steps:

1. Highlight the columns in the **Available variables** panel that you want to select as ID columns.
2. Click the right arrow to move those columns into the **Selected ID variables**.

Note: If a variable is based on a time period that is greater than your aging column, do not specify it as a class column. Instead, specify it as an ID column. For example, if your aging column is DAYDATE, do not specify WEEKDATE as another class column. Instead specify it as an ID column.

You can select the entire list of available columns by clicking the double right arrow. Back arrows can be used to deselect one, several, or all columns.

After your data has been aggregated, you can change the class list. If you remove columns from the class list, the existing data in the summarized aggregation table is merged again so that it matches the new class list. If you add new columns to the class list, then these new columns in the summarized aggregation table will have a missing value (for numeric) or blank value (for character) in the existing data. (You cannot remove a class or ID column if it is used as the source for a join column in another aggregation table. Similarly, you cannot remove a class or ID column if it is used as the source column of a rank column in this aggregation table.)

The value that is assigned to the ID variable is the last value that is read into the group that is defined by the unique combination of class variables.

Note: If an input column is given a role as a class variable, ID variable, or a statistic, it cannot simultaneously be assigned a different role. Therefore, it will not be available for selection. If all available variables are assigned as class or ID variables, then no variables will be available to be used as analysis variables. In that case, the Select analysis variables page will not be displayed. Similarly, if no analysis variables are selected, the Specify statistics page will not be displayed.

How to Specify Column Details

On the Specify column details page, you can modify the **Target Column Name**, **Target Column Label**, and **Target Column Format** for any of the input columns that you are including in the target aggregation table.

The **Target Column Name** can be any legal SAS column name. It must be unique within its table.

To modify the **Target Column Format**, double-click in the field and select the new format from the drop-down list.

How to Specify Analysis Variables

On the Select analysis variables page, you can select analysis variables that can be used to calculate statistics. If no analysis variables are selected, no statistics can be specified for this aggregation and the Specify statistics page of the wizard is not displayed.

Note: If you do not want to specify statistics, ranking, percent change, join columns, or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation table and returns you to the process flow diagram. For information, see [“Completing the Specification of the Summarized Aggregation Table” on page 157](#).

Specify Statistics

About Specifying Statistics

On the Specify statistics page, you can calculate or delete one or more statistics on analysis variables that you selected on the previous page. These variables are displayed in a grid format, which contains a row for each analysis variable. The grid contains the following columns: **Statistics**, **Weight By**, **Target Column Name**, **Target Column Description**, and **Target Column Format**. Except for the **Statistics** column, which displays the names of the selected analysis variables, the columns of the grid are blank until the statistic is specified.

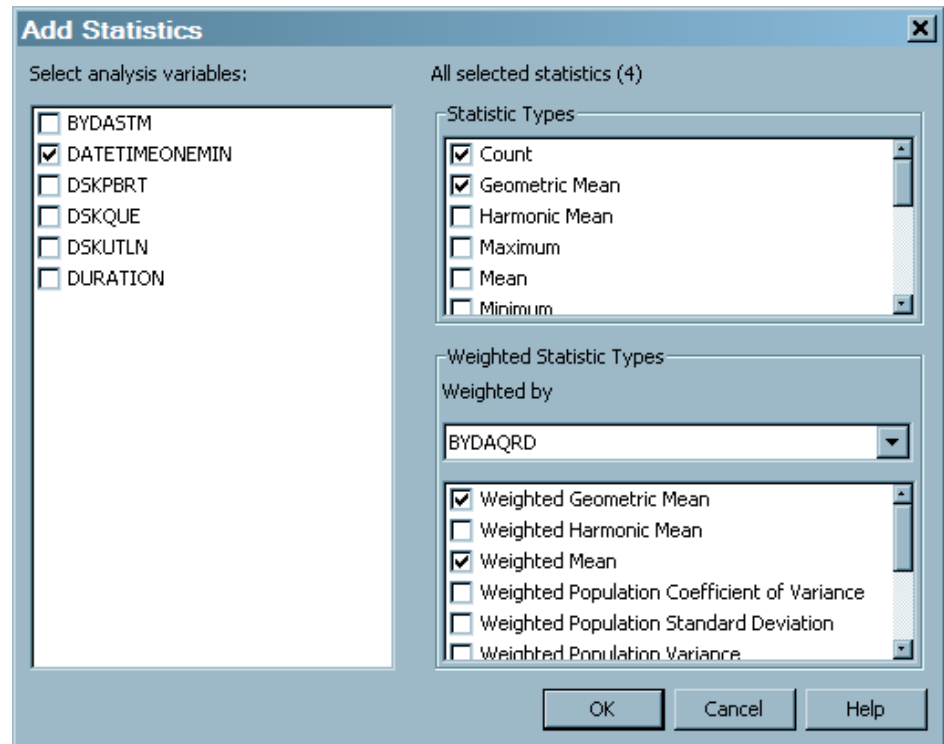
Note: For the list of statistics that SAS IT Resource Management enables you to create, see [“Statistics” on page 553](#).

Note: If you do not want to specify statistics, percent change columns, rank columns, join columns, or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation target table and returns you to the process flow diagram. For information, see [“Completing the Specification of the Summarized Aggregation Table” on page 157](#).

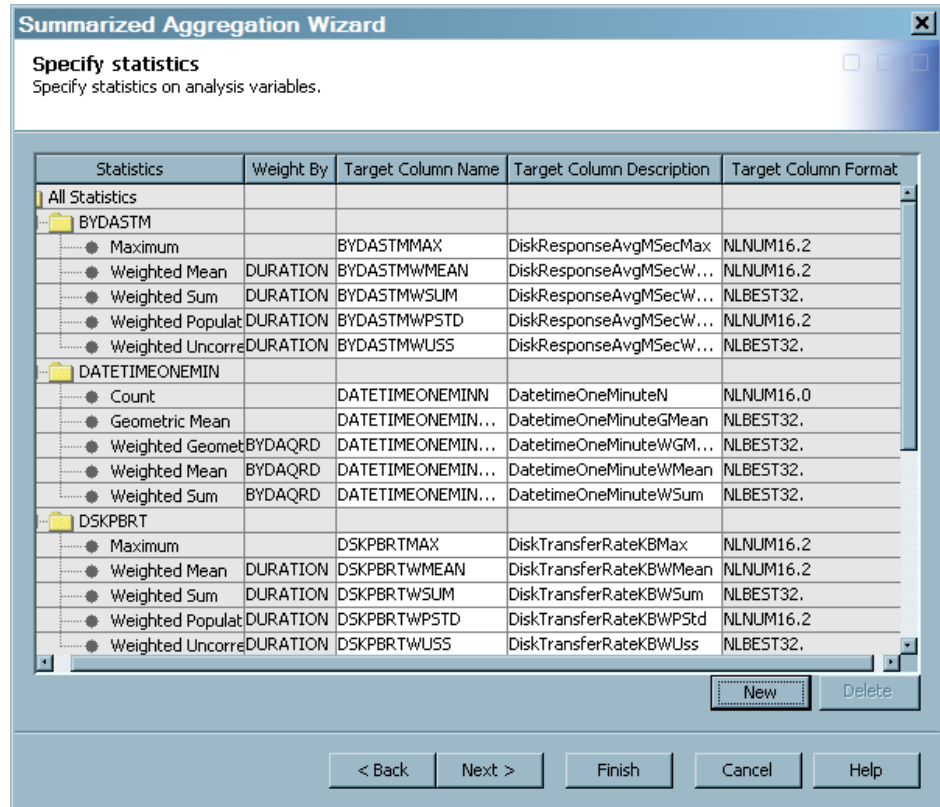
How to Create Statistics for an Analysis Variable

To create statistics for an analysis variable, perform the following steps:

1. Click **New**, which opens the Add Statistics window.
2. In the **Select analysis variables** panel, check the box next to the analysis variable for which you want to specify statistics. (Checking multiple analysis variables enables you to specify the same statistics for all the analysis variables that you selected. Otherwise, check an analysis variable and specify the statistics for it individually.)
3. To specify unweighted statistics, in the **Statistic Types** panel, check the statistics that you want to create. The counter that is next to **All Selected Statistics** displays the number of statistics that you selected.
4. To specify weighted statistics, in the **Weighted Statistic Types** panel, select the variable that you want to use from the Weighted by drop-down list. Then check the weighted statistics that you want to create.
5. In the following display, Count, Geometric Mean, Weighted Geometric Mean, and Weighted Mean statistics for the DATETIMEONEMIN analysis variable are to be created.

Display 8.9 Add Statistics Window

6. Click **OK** to add the requested and any underlying statistics to the grid under the analysis variable. The row displays these statistics along with its attributes, such as **Weight By**, **Target Column Name**, **Target Column Description**, and **Target Column Format**. These fields contain the default values that are associated with the statistics. To modify a field on this grid, highlight it, and change it as needed by typing the revision in the field.
7. The following display shows the newly added Count, Geometric Mean, Weighted Geometric Mean, and Weighted Mean statistics that will be generated for the DATETIMEONEMIN analysis variable.

Display 8.10 Specify Statistics Window with Newly Added Statistics

How to Delete a Statistic

To delete a statistic from the list of statistics that will be created for an analysis variable, perform the following steps:

1. Select the statistic.
2. Click **Delete**. If the statistic that you want to delete is used to generate another statistical column, a warning message is displayed that lists the statistical columns where the statistic is needed. You must first delete those statistical columns before you can delete this statistic. (You cannot remove a statistic column if it is used for a rank column, percent change column, or join column.)

Note: Deleting a statistic that depends on another statistic does not delete any underlying statistic.

Note: If you do not want to specify percent change, ranking, join columns, or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation table and returns you to the process flow diagram. For information, see [“Completing the Specification of the Summarized Aggregation Table”](#) on page 157.

Deleting Multiple Statistic, Percent Change, Rank, and Join Columns

In some instances, you might want to delete multiple columns at the same time. For example, you might want to delete all of the statistics for a particular analysis variable, or only some of the statistics. You can delete all of the statistics for a particular analysis variable by either selecting all the statistics individually, or by selecting the analysis variable itself. Either action will accomplish the same purpose.

However, if your selection of columns to delete includes an analysis variable and some, but not all, of the statistics for that analysis variable, a message is displayed. The message states that you selected conflicting choices and, therefore, no columns will be deleted. The message advises you to select either an analysis variable or one or more statistics.

A similar scenario can be seen in other variable selection screens, such as for percent change, rank, and join columns. In each case, you can select either the root input column or the individual output columns for deletion, but selecting a mixture of the two might be a conflict.

Specify Percent Change

About Specifying Percent Change

Percent change information is displayed in a grid format, which contains a row for each statistic that is being calculated. The grid contains the following columns: **Statistics**, **Weight By**, **Statistic Column Name**, **Target Column Name**, **Target Column Description**, and **Target Column Format**. Except for the statistic that is being calculated, all cells of the grid are initially blank.

On the Specify percent change page, you can specify the percent change for statistics. If you want to specify percent change for a variable, you must have selected an aging column. SAS IT Resource Management uses that aging column in order to calculate the period-to-period changes for that variable.

Note: If you do not want to specify percent change columns, rank columns, join columns, or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation target table and returns you to the process flow diagram. For information, see [“Completing the Specification of the Summarized Aggregation Table” on page 157](#).

How to Specify Percent Change Calculation for a Statistic

To calculate the percent change for a statistic, perform the following steps:

1. Click **New**, which opens the Add Percent Change Columns window.
2. From the list in the **Available statistics** column, select the statistics for which you want to calculate the percentage of change. Use the arrow to transfer those statistics to the **Selected statistics** column.
3. Click **OK**. The selected statistics are added to the Percent Change grid, and you are returned to the Percent Change page.

How to Remove Percent Change Calculation for a Statistic

To remove the calculation of percent change for a statistic, perform the following steps:

1. Select the statistic.
2. Click **Delete**.

Note: If the percent change column is used for a rank column in this table or a join column in another table, a dialog box is opened that displays where the column is used. You must remove those columns before you can delete the percent change column.

For information about deleting multiple percent change columns, see the information in the preceding topic called "Deleting Multiple Statistic, Percent Change, Rank, and Join Columns."

Specify Ranking

About Ranking

The ranking of data determines the position of the value of a variable or a value within a selected group of class variables. You can specify that ranks be computed on class variables, ID variables, or statistical variables. Only numeric variables can be ranked.

SAS IT Resource Management computes a dense rank, which means that the rank numbers are sequentially numbered, without gaps. The ranks start with 1 and are incremented by 1. Ties receive the same rank. For more information about ranking, see “The RANK Procedure” chapter in the *Base SAS Procedures Guide*.

Ranks are computed before the calculation of computed columns so that ranks can be used in the calculation of a computed column.

On the Specify ranking page, you can specify the variables that you want to rank and how these variables should be ranked. You can also delete a rank variable. Rank variables are displayed in a grid format, which contains a row for each defined rank variable. The grid contains the following columns: **Rank Columns**, **Target Column Name**, **Target Column Description**, **Rank Order**, **Rank Grouping**, and **Completed Days**, which are all initially blank until you add a ranking specification for a variable.

Note: Jobs that calculate rank variables might require lengthy processing time because the data might have to be read multiple times. For best results, minimize the number of rank variables that you specify.

Note: If you do not want to specify rank columns, join columns, or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation target table and returns you to the process flow diagram. For information, see “[Completing the Specification of the Summarized Aggregation Table](#)” on page 157.

How to Specify Ranking for a Variable

To specify ranking for a variable, perform the following steps:

1. Click **New** to open the Add Rank window.
2. To rank variables over all the data, click the corresponding option. (The list of variables with this option includes numeric class and ID columns that you selected from the source table as well as the statistics columns that you specified on the previous page.)

Then, in the **Select variables to rank** column, check the variables that you want to rank.

3. Alternatively, to rank variables over specific class variables, click the corresponding option. (The list of variables with this option includes only the statistics columns that you specified on the previous page.) In the **Rank over the following class variables** box, check the class variables over which you want to rank selected variables.

Then, in the **Select variables to rank** column, check the variables that you want to rank.

Note: You can select any class variables over which to rank the selected statistic. However, you should leave at least one class variable deselected. The class variables that are not selected are the variables over which the statistic is to be ranked. For example, you might have an aggregation table of average CPU Utilization, and the class variables are DAYDATE, DOMAIN, and MACHINE.

If you request a descending rank of Average CPU Utilization using a Rank Grouping of DAYDATE and DOMAIN, then that request would result in a ranking of Average CPU utilization for all machines within the DOMAIN for the given DAYDATE. Thus, the observation where the rank value is 1 would be the machine with the highest average CPU utilization, ranked separately for each unique DOMAIN and DAYDATE combination. Alternatively, if you request a descending rank of Average CPU Utilization using a Rank Grouping of DAYDATE, DOMAIN, and MACHINE, then that request would result in a ranking of Average CPU utilization for a single row of data, which is not a useful ranking.

4. In the **Rank Order** box, click **Ascending** or **Descending**.

Note: In a descending rank, the largest value gets a rank of 1, with the next largest value getting a rank of 2, and so on. In an ascending rank, the smallest value gets a rank of 1, and the next smallest value gets a rank of 2, and so on.

5. To rank only completed days, check the corresponding box.

Note: This option is available only if the **Apply completed days** option was selected on the Specify aging page of this wizard. Furthermore, if a rank is defined to have completed days, then you cannot disable the **Apply completed days** option, and you cannot disable aging.

6. Click **OK** to add the rank to the list of rank columns.

How to Delete a Rank Variable

To delete a rank variable, perform the following steps:

1. Select the variable that you want to delete.
2. Click **Delete**.

You cannot delete a rank column if it is used as the source for a join column in another aggregation table. You must first delete the join column from the other table before deleting the rank column.

For information about deleting multiple percent change columns, see the information in the preceding topic called "Deleting Multiple Statistic, Percent Change, Rank, and Join Columns."

CAUTION:

Rank columns are often used as filters in information maps. If you remove a rank column that is used as a filter in an information map, then the information map, the information map job, and the report jobs that depend on that filter might fail to run. To avoid processing errors, delete any reference to a deleted rank column in any information map, report, or transformation that uses that column.

Note: If you do not want to specify join columns or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation target table and returns you to the process flow diagram. For information, see ["Completing the Specification of the Summarized Aggregation Table"](#) on page 157.

Specify Join Columns

About Specifying Join Columns

The following rules govern the join function:

- Simple aggregations cannot serve as the source of a join column.
- Only ID, statistic, percent change, and rank columns can be joined from a summarized aggregation.
- A simple aggregation (the destination) can join columns from a summarized aggregation (the source) if and only if every class column in the source table exists as a column in the destination table.
- A summarized aggregation can join columns from another summarized aggregation if and only if the complete class list of the source table is a proper subset of the union of class and ID list of the destination table.

Note: Jobs that include joined columns might require lengthy processing time. For best results, minimize the number of joined columns that you specify.

Information about join columns is displayed in a grid format, which contains a row for each column that is to be joined. The grid contains the following columns: **Source Columns**, **Target Column Name**, **Target Column Description**, and **Target Column Format**. The entries in the grid are initially blank.

On the Specify join columns page, you can specify columns from other aggregation tables in this Aggregation transformation to join with columns in this target table.

CAUTION:

If you changed the list of class columns for an aggregation table by adding an aging column, you might have made the join invalid. This case would break the fourth rule of the rules that govern a join.

Note: If you do not want to specify join columns or a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation target table and returns you to the process flow diagram. For information, see [“Completing the Specification of the Summarized Aggregation Table” on page 157.](#)

How to Specify Join Columns

To add a join column to the table, perform the following steps:

1. Click **New**. The Add Join Columns dialog box appears.
2. Select columns from the list in the **Available Columns**. Use the arrow to transfer those columns to the **Selected Columns** panel and click **OK**. The selected columns are entered automatically into the appropriate cells of the grid, along with their corresponding default values for **Target Column Name**, **Target Column Description**, and **Target Column Format**.

How to Remove Join Columns

To delete join columns, perform the following steps:

1. Highlight the join column in the grid of columns to be joined.
2. Click **Delete**.

For information about deleting multiple percent change columns, see the information in the preceding topic called “Deleting Multiple Statistic, Percent Change, Rank, and Join Columns.”

Specify Computed Columns

About Computed Columns

A computed column stores the results of computational tasks. For example, a computed column can be used to convert unit measurements so that all the data uses the same standardized units of measurement. You can add, modify, or delete computed columns on the Add or modify computed columns page that is available in the Summarized Aggregation wizard or the Simple Aggregation wizard. By default, a computed column is numeric.

Note: If you change a computed column, the values for the computed column are recalculated for that column in all rows of the data when the job is executed. (This recalculation is performed even if no new data is added to the aggregation table.)

The Specify computed columns page displays a grid that shows the following fields for each computed column: **Name**, **Expression**, **Description**, **Length**, **Type**, **Informat**, and **Format**.

TIP A computed column called **TimePeriod** is automatically added to every simple and summarized aggregation table. Its value is always the same as the aging column, and you can refer to **TimePeriod** instead of referring to the aging column directly. This technique simplifies reporting, especially in those instances where you do not know the name of the aging column. (If no aging is specified for the target table, then the field is set to missing.) The expression for this column is defined when the aggregation wizard is completed. For example, if the aging column of an aggregation table is DAYDATE, the expression for the TimePeriod computed column is DAYDATE and its expression is **TIMEPERIOD=DAYDATE;**

On the Specify computed columns page, you can create, modify, or delete computed columns.

How to Add a Computed Column

To add a computed column to the aggregation table, on the Computed Columns page of the aggregation wizard, perform the following steps:

1. Click **New**. A new line is displayed on the grid that contains default values for the following attributes: name, length, type, informat, and format. (The default value for the name of the column is “Untitled n ”, where n is the n th computed column for this aggregation.)

Display 8.11 Specify Computed Columns

Summarized Aggregation Wizard

Specify computed columns
Add or modify computed columns.

#	Name	Expression	Description
1	TimePeriod	monthDate	TimePeriod
2	ipIncomingAddrErrorRate	ipIncomingAddrErrorCountS...	IpIncomingAddress...
3	ipIncomingDiscardRate	ipIncomingDiscardCountSU...	IpIncomingDiscard...
4	ipIncomingHdrErrorRate	ipIncomingHdrErrorCountS...	IpIncomingHeader...
5	ipIncomingNoRouteRate	ipIncomingNoRouteCountS...	IpIncomingNoRout...
6	ipIncomingReceiveRate	ipIncomingReceiveCountSU...	IpIncomingReceive...
7	ipIncomingTruncPktRate	ipIncomingTruncPktCountS...	IpIncomingTruncD...
8	ipIncomingUnkProtoRate	ipIncomingUnkProtoCountS...	IpIncomingUnkProt...
9	ipOutgoingDiscardRate	ipOutgoingDiscardCountSU...	IpOutgoingDiscard...
10	ipOutgoingNoRouteRate	ipOutgoingNoRouteCountS...	IpOutgoingNoRout...
11	ipOutgoingRequestRate	ipOutgoingRequestCountS...	IpOutgoingReques...
12	Untitled12		

New Delete

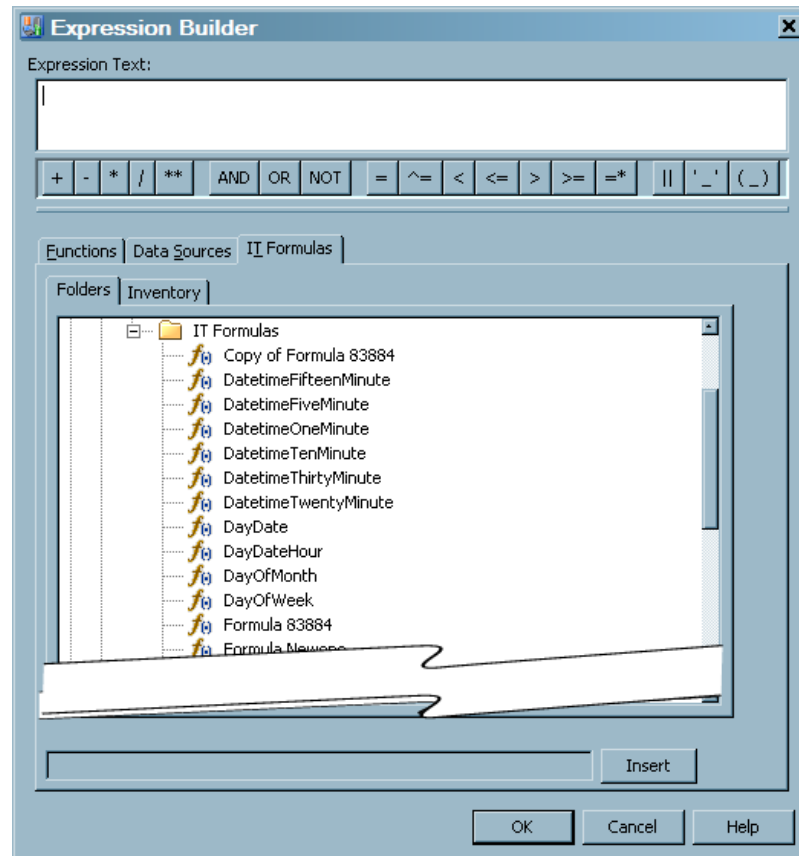
< Back Next > Finish Cancel Help

2. Enter the values that you want to use for the **Name**, **Expression**, and **Description** fields. You can enter values for the **Length**, **Type**, **Informat**, and **Format** fields, or you can accept the default values for these fields.

How to Modify a Computed Column

To modify a computed column, perform the following steps:

1. Select the column from the drop-down list. You can change the values of this column as needed.
2. In the **Name** field, you can change the name of your computed column. The name of the computed column can be any legal SAS name. It must be unique within the target table.
3. In the **Expression** field, you can specify a modified computation for this computed column by entering it in this field or by using the Expression Builder function that is provided by SAS Data Integration SAS IT Resource Management. Double-click in the field and click ... to open the Expression Builder window.

Display 8.12 Expression Builder Dialog Box with IT Formulas Tab

Apart from the many functions that are available from the **Functions** and **Data Sources** tabs of this window, the Expression Builder window provides another tab to assist you in defining the computation for this column—the **IT Formulas** tab.

- The **Functions** tab of the Expression Builder window lists a variety of functions, conversions, and other actions that you can specify in your expression. Select the function that you want to use in the expression and click **Insert**.
- The **Data Sources** tab of the Expression Builder window lists all the columns that are generated in the aggregation table output. Select the column that you want to use in the expression and click **Insert**.

For simple aggregations, the **Data Source** tab displays all the output columns that you selected for the simple aggregation table. For the summarized aggregation table, the **Data Source** tab displays all the class, ID, statistic, percent change, rank, and join columns that you selected for the summarized aggregation table. In addition, it displays the system-generated LastUpdated, CompletedDay, and ContribCount variables.

All variable names that are used in the computation are the output column names. For example, if you choose to use an input column named SYSTEM as a class variable, you might rename it MACHINE. Then, if you want to use that column when defining a computed column, you should refer to that column as MACHINE, not SYSTEM.

- The **IT Formulas** tab lists the formulas that you have defined as well as the formulas that are supplied with SAS IT Resource Management. Select the formula that you want to use and click **Insert**.

In addition to the conventional rValue expression (where the expression consists of code that is appropriate only for the right-hand side of an assignment statement), SAS IT Resource Management also supports more complex expressions. For example, your expression can use SAS code that might include loops, IF statements, and so on. This code must be written in valid SAS DATA step syntax.

When you are satisfied with the expression that you defined, click **OK** in the Expression Builder window to close that window and place the expression in the appropriate field of the computed column.

Note: If invoked from the **Filters** tab of an Aggregation transformation, the Expression Builder window does not contain the **IT Formulas** tab. This tab is available only when the window is invoked from the **Add or modify a computed column** page of the wizards for summarized and simple aggregations.

For more information about the Expression Builder window, click the F1 key or **Help** from within that window.

4. In the **Description** field, you can change the descriptive information about the computed column.
5. In the **Type** field, you can change the type of the computed column. Double-click in the field and use the arrow to display the drop-down list of valid types (either character or numeric) from which you can select the type for this computed column. Fields that are the result of a calculation should be specified as numeric. By default, the computed column is numeric.
6. In the **Length** field, you can change the length of the computed column. For numeric type columns (not character type columns), this numeric value must be from 2 through 8. (Reducing the length of a numeric computed column might introduce precision errors.) For character computed columns, the length can be from 1 through 32767.
7. The **Informat** field is not used by the Aggregation transformation.
8. In the **Format** field, you can change the SAS format of the computed column. Double-click in the field and from the drop-down list of valid formats, select the format for this computed column.

You can also enter a format name if it will be available when you execute the job or when you view the data in SAS.

How to Delete a Computed Column

Deleting a computed column removes it from the metadata for a table. The column will no longer be generated when the Aggregation transformation job is run. However, deleting a computed column does not remove the column from the physical tables of aggregated data that have already been generated.

To delete a computed column from an aggregation table, perform the following steps:

1. Select the computed column that you want to delete.
2. Click **Delete**.

When you have finished working with your computed columns, click **Next** to continue to the final page of the wizard.

Note: After the job is redeployed and executed, changes to the computed columns will be reflected in the physical table. For information about redeploying jobs, see [“Redeploy All Jobs on the Server” on page 281](#).

For information about deleting multiple percent change columns, see the information in the preceding topic called "Deleting Multiple Statistic, Percent Change, Rank, and Join Columns."

Completing the Specification of the Summarized Aggregation Table

The final page of the wizard displays the details for the summarized aggregation table that you specified. If you are satisfied with your choices, click **Finish**. Click **OK** to return to the process flow diagram.

Aggregation tables are stored in sorted order, with an asserted **SORTEDBY=** data set option. The SORTEDBY assertion reflects that the data is stored in the following order:

1. aging column
2. alphabetical list of the remaining class columns

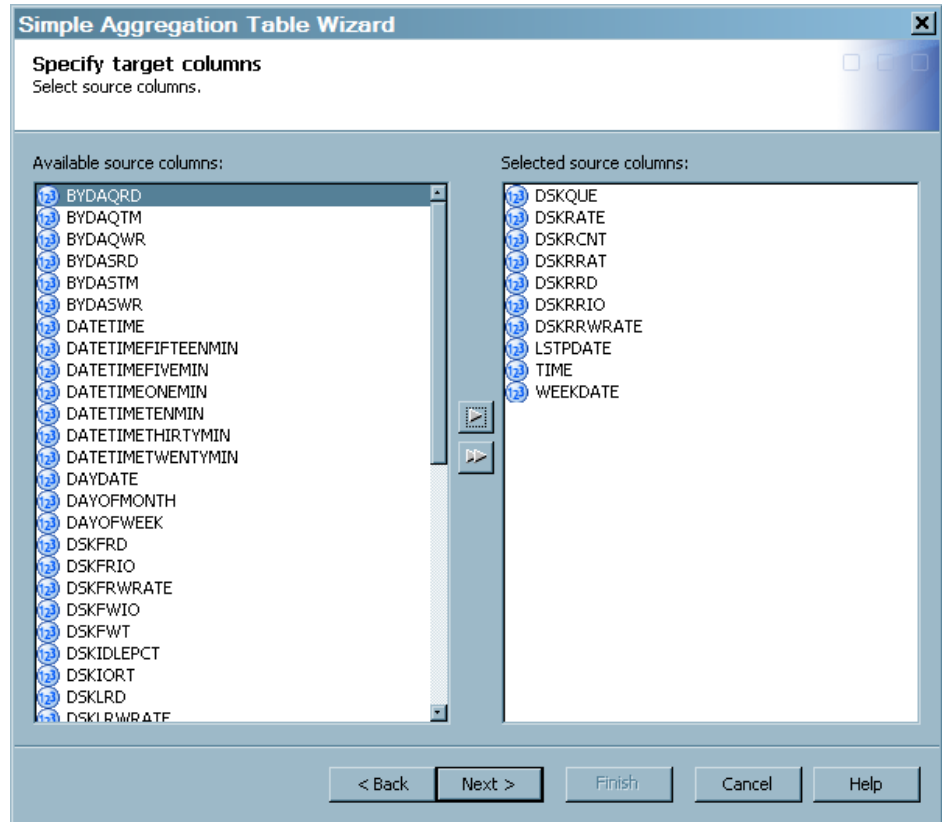
BY-processing can be performed using the class list, if the BY list matches the same SORTEDBY assertion. With SAS IT Resource Management, you can leverage SAS Data Integration Studio to define additional indexes for use where BY-processing is required. If you require a sorted aggregation table, you can input that table to a SORT transformation and store the resulting table wherever you want.

For information about indexes, see ["Indexing an Aggregation Table" on page 173](#).

Creating an Aggregation Table with the Simple Aggregation Table Wizard

To create an aggregation table with the Simple Aggregation Table wizard, respond to the wizard's prompts.

1. To enter general information about the table, see ["How to Enter General Information about the Aggregation Table" on page 141](#).
2. If this aggregation table is the first target table that is specified for this Aggregation transformation, an additional page of the wizard opens. This page prompts you for the physical target location and for information about the SAS library where the data will be stored. For information about specifying a library, see ["How to Specify Library Information" on page 141](#).
3. Filters can limit the amount of data that is input to an aggregation table. On the page of the wizard that is called Select filter and column for data duplication, you can specify a filter. For information, see ["How to Select a Filter" on page 142](#).
4. On the Specify aging page, you can specify aging and purging criteria. For information, see ["Specify Aging" on page 142](#).
5. On the Specify target columns page, you can specify the columns from the source table that are to be included in the target aggregation table. To select columns for the target table, highlight the columns in the **Available source columns** panel and then click the right arrow to move those columns into the **Selected source columns**.

Display 8.13 Specify Target Columns Page of the Simple Aggregation Table Wizard

You can select the entire list of available source columns by clicking the double right arrow. Back arrows can be used to deselect one, several, or all columns.

If aging is to be performed for this simple aggregation table, then the selected aging column is automatically added to this list. It cannot be deleted while it is still in use as the aging column.

6. On the Specify column details page, you can modify the **Target Column Name** and **Target Column Label** fields for any of the source columns that you are including in the target aggregation table. To do so, enter the new value in the field.

To modify the format, double-click in the field and select the new format from the drop-down list.

7. On the Join columns page, you can specify columns from other aggregation tables in this Aggregation transformation to join with columns in this target table. For information, see [“Specify Join Columns” on page 151](#).

Note: The Join Columns page is available from the Simple Aggregation wizard only if there are columns from separate target tables that are eligible for joining.

Note: If you do not want to specify a computed column for this aggregation table, click **Finish**, which updates the Aggregation transformation with this new aggregation target table and returns you to the process flow diagram.

8. On the Computed columns page, you can create, modify, or delete computed columns.

For information about how to add, modify, or delete computed columns, see [“Specify Computed Columns” on page 153](#).

9. The final page of the wizard displays the details for the simple aggregation table that you specified. If you are satisfied with your choices, click **Finish**. Click **OK** to return to the process flow diagram.

Creating Aggregation Tables with the Aggregations through Templates Wizard

Aggregation tables can be created using those tables that are supplied by the SAS IT Resource Management Solution as a starting point. To create an aggregation table with the New Aggregation Table from Templates wizard, respond to the wizard's prompts.

1. On the Specify location page, you can specify the metadata location where the newly created tables will be saved. You must specify a location for the staged table that is within the IT data mart that is specified in the **IT Data Mart** field.

You can override the default value for this field by clicking **Browse** and selecting a path to the table from the Select a Location dialog box that displays the metadata directory structure.

Note: The **IT data mart** field displays the name of the IT data mart where the job is located. It cannot be modified.

2. On the Select templates page, check the template tables that you want to use to create aggregation tables. The **Total number of selected template tables** field is updated to reflect the number of template tables that you selected.

For staged tables that are designed to contain data that is already summarized to an hourly interval, you cannot successfully select aggregations that attempt to summarize the data to smaller intervals. When you specify an aggregation table that summarizes data for a time period that is shorter than the source table normally provides, a warning message is displayed stating that the aging column cannot be created and the aggregation table is not created. To avoid this situation, make sure that your source table provides the appropriate level of summarized values for the aggregation that you want to generate. For example, if you want to specify an aggregation table such as FifteenMinAcctTcodeSy, make sure that the source data (that is, the staged table) is summarized no higher than 15 minutes and import the DATETIMEFIFTEENMINUTE computed column into the staged table. (For information about importing a column into a staged table, see [“Modify a Staged Table” on page 122.](#))

Note: If the templates that you select have a dependency on other templates, an information box is displayed that lists the additional aggregations that will be created. Click **Yes** to accept this action and to continue. Click **No** to return to the Select Templates page.

3. If this aggregation table is the first target table that is specified for this Aggregation transformation, an additional page of the wizard opens. This page prompts you for the physical target location and for information about the SAS library where the data will be stored. When this page opens, the following fields are already specified with default values. (However, you can modify these fields.)
 - The **Name** field contains the name of the library that will contain all aggregation tables that are associated with the Aggregation transformation that you are working with. This field can contain up to 60 characters and should be unique within the IT data mart. A value is required in this field.

- The **Description** field contains the information that describes the library. It can contain no more than 200 characters. It is an optional field.
- The **Server** field contains the application server that you specified for the IT data mart that you are working with. The default application server is the SASITRM server. A value is required in this field.
- The **Libref** field contains the name that is temporarily associated with the library. The libref for an aggregation library is a concatenation of AGG and a random number (for example, *AGG27040*). A value is required in this field.
- The **Path specification** field contains the location where the tables for this library are to be stored. You must specify a location for the tables that is within the IT data mart that is specified in the **IT Data Mart** field.

Enter this field manually or click **Browse** to select a location from the Select a Directory dialog box that displays the directory structure on the server's file system. A value is required in this field.

4. The final page of the wizard displays the details for the aggregation tables that you specified. If the aggregation job that you specified adds more than 20 tables to an existing job, a warning message is displayed. The message suggests that you might want to remove some tables in order to experience better performance. Click **Yes** to continue to generate the tables. Click **No** to return to the Select templates page, where you can change your selections. If you are satisfied with your choices, click **Finish**.

Click **OK** to return to the process flow diagram.

Note: Some aggregation templates require other aggregation tables to be created. If so, a dialog box appears that notes the dependency and asking for confirmation before proceeding.

Modifying Aggregation Tables

Delete an Aggregation Table from an Aggregation Transformation

To delete an aggregation table from an Aggregation transformation, perform the following steps:

1. Locate the Aggregation transformation that generates the aggregation table that you want to delete. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the aggregation table that you want to delete.
3. From the drop-down list, select **Delete**. The table is removed from the Aggregation transformation that is displayed in the process flow diagram. However, the metadata for that table and the library are not deleted from the IT data mart. For information about this topic, see [“Delete an Aggregation Table from the IT Data Mart”](#) on page 161.

In order for your changes to take effect, you must redeploy the job. For information about redeploying jobs, see [“Redeploy All Jobs on the Server”](#) on page 281.

CAUTION:

To avoid processing errors, delete any reference to this table, as a source or target for any transformation in any other job. For example, Information Map transformations that use the deleted table as a source table should be deleted.

Delete an Aggregation Table from the IT Data Mart

To delete an aggregation table from an IT data mart, perform the following steps:

1. Locate the aggregation table that you want to delete. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the aggregation table is generated.
2. Expand the IT data mart tree and navigate to the aggregation table that you want to delete. Right-click the table.
3. From the drop-down list, select **Delete**.
4. Click **Yes** to confirm that you want to delete this table. The table is removed from the IT data mart. However, the associated library is not deleted.

CAUTION:

Deleting a table can break join columns if the table that you delete is the source for another table.

Delete an Aggregation Table with Contents

To delete with contents an aggregation table from an IT data mart, perform the following steps:

1. Locate the aggregation table that you want to delete. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the aggregation table is generated.
2. Right-click the aggregation table that you want to delete.
3. From the drop-down list, select **Delete with Contents**.
4. Click **Yes** to confirm that you want to delete this table. The table is removed from the IT data mart. However, the associated library is not deleted.

CAUTION:

Deleting a table can break join columns if the table that you delete is the source for another table.

Erase an Aggregation Table

Erasing an aggregation table deletes the physical table, its contents, and its metadata. To erase an aggregation table, perform the following steps:

1. In a tree view of SAS IT Resource Management, locate the aggregation table that you want to erase.
2. Right-click the aggregation table that you want to erase. The aggregation job that includes the aggregation table must be closed and currently not in use before erasing.
3. Select **Erase**.
4. Click **Yes** to confirm that you want to erase this aggregation table.

CAUTION:

If you erase an aggregation table that serves as a source table in other jobs, then these associated jobs might not execute because they are missing a source table. Simply erasing an aggregation table does not affect the metadata of the associated aggregation or information map jobs. However, if you erase an aggregation table and then rerun the aggregation job, then the data for the deleted aggregation table is not created and the jobs that used that table will fail due to a missing source table.

Modify a Simple or a Summarized Aggregation Table

If you need to modify an aggregation table, for best results do so from within the aggregation job. (If you want to edit an aggregation table that was created by the Simple Aggregation Table wizard or the Summarized Aggregation Table wizard, you must first save the job that created that aggregation table.) To modify a simple or a summarized aggregation table, perform the following steps:

1. Locate the Aggregation transformation that generates the aggregation table that you want to modify. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the aggregation table that you want to modify. From the list that displays, select **Edit table type Aggregation Table**.

Note: If the table is a simple aggregation table, the list displays **Edit Simple Aggregation Table**. If the table is a summarized aggregation table, the list displays **Edit Summarized Aggregation Table**. However, if the aggregation table is not associated with an Aggregation transformation, then the **Edit table type Aggregation Table** option will not be available.

3. Use the **Jump to Page** option to go directly to the page of the aggregation wizard that you want to change. Use the drop-down list to display the pages to which you can go directly.

Display 8.14 Jump to Percent Change Page for a Summarized Aggregation Table

4. Select the page that contains the specifications that you want to modify and click **Next**.
5. Modify the contents of the page that you selected. (For information about which fields can be changed, follow the instructions that apply to the creation of the type of table that you want to modify.)
 - For information about creating summarized aggregation tables, see [“Creating Aggregation Tables with the Summarized Aggregation Table Wizard”](#) on page 141.
 - For information about creating simple aggregation tables, see [“Creating an Aggregation Table with the Simple Aggregation Table Wizard”](#) on page 157.

Note: If you try to delete a column that is used to create a statistic, percent change, rank, or join column, a message is displayed that lists the column, where it is used, and its table name. If a column was deleted by means of the **Columns** tab of the Properties dialog box, then the next time the aggregation wizard is opened on that table, a message is displayed that indicates any columns whose source column is no longer available. Click **Yes** to delete the columns from the table and continue to edit the table. Click **No** to keep the columns. In that case, the wizard is not opened.

6. Click through the remaining pages of the wizard until the summary page appears.
7. Click **Finish** to save your changes and return to the process flow diagram.

In order for your changes take effect, you must redeploy the job. For information about redeploying jobs, see [“Redeploy All Jobs on the Server”](#) on page 281.

Purging the Content of an Aggregation Table

Purge Data from an Aggregation Table from the Process Flow Diagram

To purge data from an aggregation table, perform the following steps:

1. Locate the Aggregation transformation that generates the aggregation table that you want to modify. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the aggregation table whose data you want to purge.
3. From the drop-down list, select **Purge**.
4. Click **Yes** to confirm that you want to purge the data from this aggregation table.

Purge Data from an Aggregation Table on the IT Data Marts Tab

You can also purge the contents of an aggregation table by performing the following steps:

1. Locate the Aggregation transformation that generates the aggregation table that you want to modify. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located.
2. Right-click the aggregation table whose data you want to purge.
3. From the drop-down list, select **Purge**.
4. Click **Yes** to confirm that you want to purge the data from this aggregation table.

Renaming an Aggregation Table

CAUTION:

You can change the name of the physical aggregation table only by using the Edit Aggregation wizard. Changing the name of the table on the **General** tab of the Properties dialog box does not change the name of the physical table.

To rename a simple or a summarized aggregation table, perform the following steps:

1. Locate the Aggregation transformation that generates the aggregation table that you want to modify. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the aggregation table that you want to rename.
3. From the drop-down list, select **Edit <aggregation-type> Aggregation Table**, where *aggregation-type* is Summarized or Simple. This action opens the Summarized Aggregation wizard or the Simple Aggregation wizard.
4. To change the name of the table in the metadata, enter the new name of the aggregation table in the **Name** field on the first page of the wizard.

To change the name of the physical table, enter the new name of the aggregation table in the **SAS Table Name** field.

5. Click **Next** to proceed through all the pages of the wizard. Then click **Finish**.

6. Click **OK** to save your changes and return to the process flow diagram.

Note: This task changes the name of the table in the metadata. The name of the physical SAS data set is changed when you click **Finish**. In order for your changes to take effect, you must redeploy the aggregation job. In addition, even though the metadata linkage to that renamed table still exists, it is advisable to redeploy the information map job that uses the renamed table as input. For more information about redeploying jobs, see [“Redeploy All Jobs on the Server” on page 281](#).

Working with Aggregation Transformation Filters

About Aggregation Transformation Filters

A filter can be used to limit the amount of data that is input to an individual aggregation table. You can define a filter to be used in an Aggregation transformation, modify or edit the filter, and delete it from an Aggregation transformation. A filter that is specified for an Aggregation transformation can be applied to any of the aggregation tables that are being specified in the Aggregation transformation. However, the filter will not affect any of the data that is already in the tables that were created or updated before the application of the filter. If you define a filter to be used in an Aggregation transformation, modify an existing filter in an Aggregation transformation, or delete a filter from an Aggregation transformation, the existing aggregation tables are not automatically regenerated.

Note: Filters are used only when data is read into the aggregation table. If you modify a filter, the data that is already in the table remains unchanged.

After a filter has been added to an Aggregation transformation, it can be applied to the input of any aggregation table within that transformation. For both the Simple Aggregation Table wizard or the Summarized Aggregation Table wizard, you can select an input filter on the Select filter and column for data duplication page of that wizard. When you are using the Aggregations through Templates wizard, if the aggregation that you select uses a filter, it will be added to the specification automatically.

Access the ITRM Options Tab on the Properties Dialog Box of an Aggregation Transformation

To define, modify, or delete a filter for an Aggregation transformation, you must first access the **ITRM Options** tab on the Properties dialog box of that Aggregation transformation. To do so, perform the following steps:

1. From the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor window, and locate the appropriate Aggregation transformation.
3. Right-click the Aggregation transformation. From the drop-down list, select **Properties**.
4. Select the **ITRM Options** tab to open the grid of input filters. The **List of Filters** group box displays the name and expression of input filters that are defined for this Aggregation transformation.

Define a Filter for an Aggregation Transformation

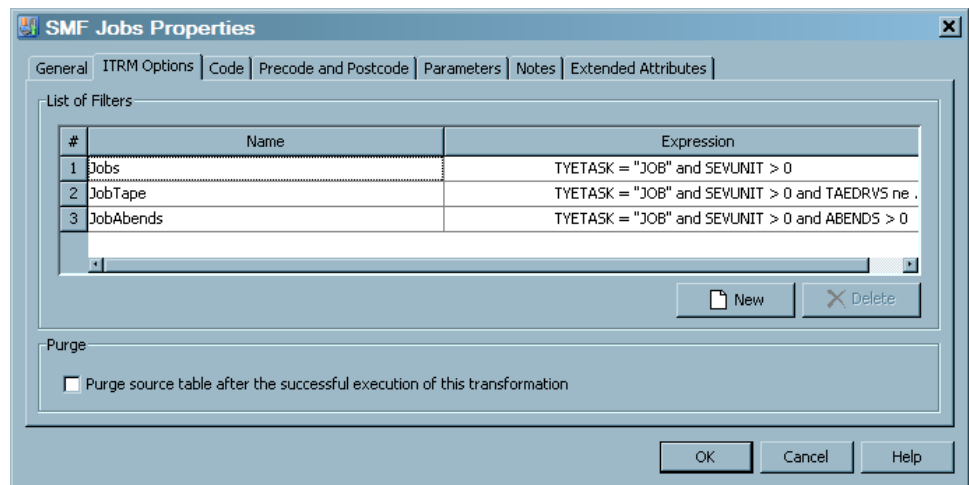
To define an input filter for an Aggregation transformation, perform the following steps:

Access the **ITRM Options** tab on the Properties dialog box of the Aggregation transformation for which you want to define a filter. To do so, see the preceding topic.

1. Click **New** to display a new line in the grid on which you can enter a new filter. The default name of this filter is "Untitled n ," where n is the n th filter that exists for this Aggregation transformation.

The following display of the **ITRM Options** tab shows that a new filter is being defined for the SMF Jobs Aggregation.

Display 8.15 List of Filters Group Box



2. Enter the name of the filter. Consider giving the filter a name that is unique within this Aggregation transformation.
3. Enter the expression that specifies the filter. You can enter the expression directly in this field, from a formula, or by using the Expression Builder functionality of SAS Data Integration SAS IT Resource Management. For information about how to use this functionality, see [“About the Expression Builder” on page 168](#).
4. If you want to purge the source table upon successful creation of aggregation tables, click the corresponding check box.
5. When you are satisfied with the definition for the filter, click **OK** to save your work and return to the process flow diagram.

Delete a Filter from an Aggregation Transformation

To remove a filter from an Aggregation transformation, you must first remove the filter from all the aggregation tables that use it. (For information about how to do so, see [“Delete a Filter from an Aggregation Table” on page 167](#).)

Access the **ITRM Options** tab on the Properties dialog box of the Aggregation transformation from which you want to delete a filter. To do so, see [“Access the ITRM Options Tab on the Properties Dialog Box of an Aggregation Transformation” on page 165](#). Then perform the following steps:

1. Select the filter that you want to delete.

2. Click **Delete**. If the filter that you selected is used in another table, a warning message is displayed that lists the tables where the filter is used. In that case, you must delete the filter from those tables.
3. Click **OK** to save your changes and return to the process flow diagram.

Note: In order to apply the changes that you made, you must redeploy the jobs that use that filter. For information about deploying jobs, see [“Redeploy All Jobs on the Server” on page 281](#).

Delete a Filter from an Aggregation Table

An aggregation filter that is being used by an aggregation table cannot be deleted from an Aggregation transformation. If you attempt this action, a message box appears that identifies the aggregation tables that are using the filter. To delete the input filter from the Aggregation transformation, you must first remove it from all aggregations where it is being used.

To remove a filter from an aggregation table, perform the following steps:

1. From the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the aggregation table that uses the input filter. From the drop-down list, select **Edit table type Aggregation Table**.

Note: If the table is a simple aggregation table, the drop-down list displays **Edit Simple Aggregation Table**. If the table is a summarized aggregation table, the drop-down list displays **Edit Summarized Aggregation Table**.

3. Use the **Jump to Page** option to go directly to that page of the wizard. The drop-down list shows the pages that you can go to. Select **Apply Filter** and click **Next**.
4. In the **Select filter** field, the drop-down list shows the filters that can be applied. You can choose another filter to apply to the table, or else the **(None)** option, which removes the filter from this table.

Note: Only one filter can be applied to an aggregation table.

5. Click through the remaining pages of the wizard until the summary page displays.
6. Click **Finish** to save your changes and return to the process flow diagram.

Modify a Filter of an Aggregation Transformation

To modify or edit a filter that is associated with an Aggregation transformation, perform the following steps:

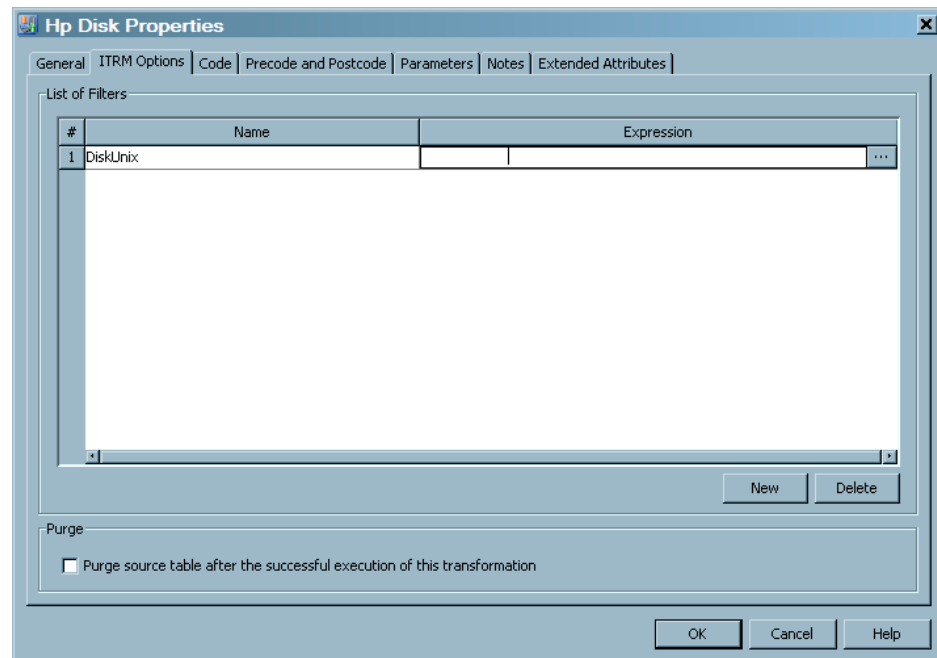
1. Access the **ITRM Options** tab on the Properties dialog box of the Aggregation transformation for which you want to modify a filter. To do so, see [“Access the ITRM Options Tab on the Properties Dialog Box of an Aggregation Transformation” on page 165](#).
2. Select the filter that you want to modify. You can change the name and the expression of the filter. You can also change the specification about purging the input table after the aggregation tables have been created.
 - To change the name of the filter, select the **Name** field of the filter to highlight it. Then enter the new name of the filter.

- To change the expression, select the **Expression** field to highlight it. Then enter the modified expression that specifies the filter. You can enter the expression directly in this field or you can use the Expression Builder functionality of SAS Data Integration SAS IT Resource Management. For information about how to work with the Expression Builder in SAS IT Resource Management, see [“About the Expression Builder” on page 168](#).
 - To change the specification about purging the input table after the aggregation tables are created, click the corresponding check box. (A checked box means that the source table will be purged.)
3. When you have finished modifying the filters, click **OK** to save your changes and return to the process flow diagram.

About the Expression Builder


To invoke the Expression Builder, double-click the **Expression** field.

Display 8.16 Invoking the Expression Builder



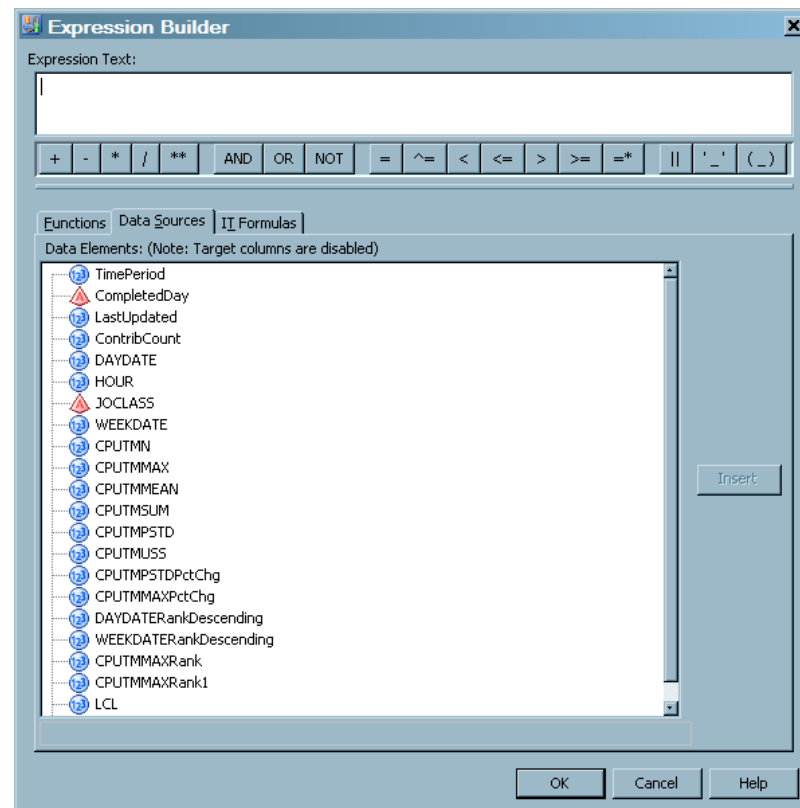
Then, double-click the icon that contains the ellipsis (...). The Expression Builder window appears and enables you to specify a filter.

Note: When invoked from the Add Filter dialog box, the Expression Builder window consists of two tabs: **Functions** and **Data Sources**.

Display 8.17 Expression Builder Window with Functions Tab Selected

The **Data Sources** tab of the Expression Builder window lists the columns of the table that is the input to this Aggregation transformation.

Display 8.18 Expression Builder Window with Data Sources Tab Selected



You can select any of these columns from the list of data sources to be an operand in the expression that you are building.

For information about the Expression Builder window, click the F1 key or **Help** from within that window.

Note: The syntax of the expression that specifies the filter must conform to the rules for an IF statement and should result in a nonzero value for true and zero value for false. (Do not code the word **IF**; code the expression only.) Certain operations that are valid only in a WHERE clause fail if they are used in a filter. For information about the IF and WHERE statements, see the *SAS Statements: Reference*.

When you are satisfied with the expression, click **OK** in the Expression Builder window to close that window and place the expression in the appropriate field of the filter.

Working with Aggregation Transformations

Properties of an Aggregation Transformation

To access an Aggregation transformation, from the IT Data Marts tree navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window. Right-click the Aggregation transformation and select **Properties**.

The Properties dialog box appears. It consists of the following tabs:

- The **General** tab contains the identifying information about the transformation. The fields on this tab can be modified.

Display 8.19 General Tab of an Aggregation Transformation

SMF Jobs Properties

General | ITRM Options | Code | Precode and Postcode | Parameters | Notes | Extended Attributes

Name: SMF Jobs

Description:

OK Cancel Help

The following display shows the **General** tab of an aggregation table.

For items in the preceding display, the following definitions apply:

Name

identifies the name of the Aggregation transformation.

Description

describes the Aggregation transformation.

- The **ITRM Options** tab enables you to view, add, modify, and delete filters for an Aggregation transformation. It also enables you to specify that the Aggregation transformation should purge data from the source table after the Aggregation transformation has executed successfully.

Display 8.20 ITRM Options Tab of an Aggregation Transformation

SMF Jobs Properties

General | ITRM Options | Code | Precode and Postcode | Parameters | Notes | Extended Attributes

List of Filters

#	Name	Expression
1	Jobs	TYETASK = "JOB" and SEVUNIT > 0
2	JobTape	TYETASK = "JOB" and SEVUNIT > 0 and TAEDRV5 ne
3	JobAbends	TYETASK = "JOB" and SEVUNIT > 0 and ABENDS > 0 ...

New Delete

Purge

☐ Purge source table after the successful execution of this transformation

OK Cancel Help

Name

The **Name** column displays the list of filters that are available for an Aggregation transformation. (You can apply only zero or one filter to an aggregation target table.) A value is required in this field.

Expression

The **Expression** field displays the definition or the calculation of the filter. A value is required in this field.

New

Clicking the **New** option adds a new row in the table that displays all the available filters. This action enables you to specify a filter that can be applied to a table that is input to an Aggregation transformation.

Delete

Clicking the **Delete** option deletes the selected filter.

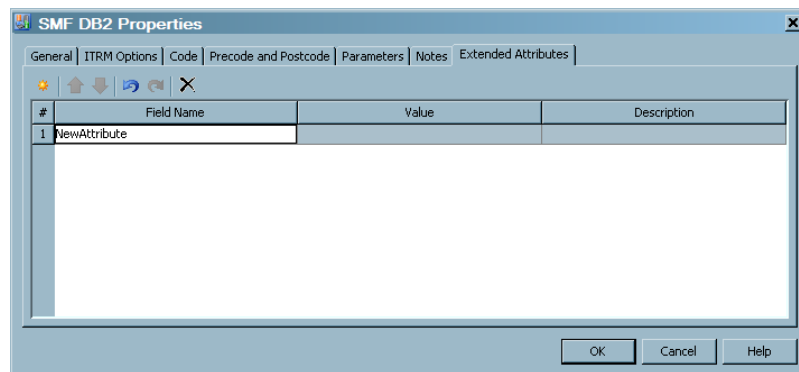
Purge source table after the successful execution of this transformation

Checking the **Purge source table after the successful execution of this transformation** option removes the data from the source table after the aggregation tables are processed.

- The **Code** tab contains the code that is generated for the transformation.
- The **Precode and Postcode** tab contains information about user-written code that might be inserted at the beginning or the end of the transformation.
- The **Parameters** tab contains information about prompts.
- The **Notes** tab contains areas where you can annotate the table.
- The **Extended Attributes** tab contains a custom property that is not part of the metadata for the transformation.

The following display shows the **Extended Attributes** tab of an aggregation table.

Display 8.21 Extended Attributes Tab of an Aggregation Transformation



For items in the preceding display, the following definitions apply:



adds a row to the attribute table. A number and a default name are provided, but you must enter the value and description.



if present, enables you to reorder rows by moving a selected row up in the table.



if present, enables you to reorder rows by moving a selected row down in the table.



deletes a selected attribute from the list.

#	specifies the number of the extended attribute.
Field Name	specifies the attribute name to be defined.
Value	specifies the value for the attribute that is specified in the Field Name column.
Description	describes the attribute and value.

Delete an Aggregation Transformation from a Job

A job can consist of multiple transformations. An Aggregation transformation can consist of multiple aggregations, each generating an aggregation table. To delete an Aggregation transformation from a job, perform the following steps:

1. Locate the Aggregation transformation that you want to delete. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the box that represents the Aggregation transformation that you want to delete.
3. From the drop-down menu, select **Delete**.

Note: If you delete an Aggregation transformation, the target tables that were populated by those aggregations are not deleted. However, the tables are no longer updated.

Rename an Aggregation Transformation

1. Locate the Aggregation transformation that you want to rename. To do so, from the IT Data Marts tree, navigate to the IT data mart that contains the job where the Aggregation transformation is located. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the Aggregation transformation box.
3. From the drop-down menu, select **Properties**.
4. On the **General** tab, enter the new name of the Aggregation transformation.
5. Click **OK** to save your changes and return to the process flow diagram.

Note: Renaming an Aggregation transformation does not rename the aggregation tables that it generates.

Indexing an Aggregation Table

Introduction to Indexes

An index is a component of a SAS data set that enables SAS to access observations in the SAS data set quickly and efficiently. The purpose of SAS indexes is to optimize

WHERE clause processing and to facilitate BY-group processing. SAS IT Resource Management supports indexes for summarized and simple aggregation tables.

- Summarized Aggregations: You can create additional indexes manually.

CAUTION:

Do not create unique indexes for summarized aggregation tables.

- Simple Aggregations: Indexes are not automatically created for input columns that are specified for aggregation tables. However, you can create indexes manually.

For more information about indexes, see *SAS Data Integration Studio: User's Guide* or the Help for SAS Data Integration Studio.

How to Access an Index

Indexes are physically created when the aggregation job is run. To access an index, perform the following steps:

1. Right-click the aggregation table and select **Properties** from the drop-down list.
2. Select the **Indexes** tab.

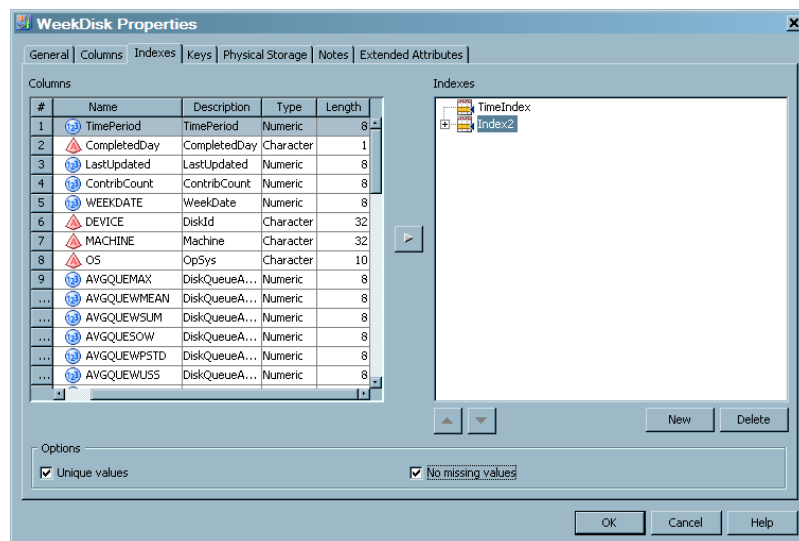
The left panel of the **Indexes** tab contains a list of the variables that can be used to index the data.

Note: Any variables that are in the table can also be in the index. For a simple aggregation table, these variables were previously specified as input columns for the table. For a summarized aggregation table, these variables were previously specified as class, ID, statistic, rank, percent change, join, or computed columns for the table.

The right panel contains the indexes that are specified for this aggregation.

In the following display, an index called TimeIndex has already been created. A new index called Index2 is in the process of being created.

Display 8.22 Indexes Tab of an Aggregation Table



From the **Indexes** tab, you can create, rename, and delete indexes.

Create an Index

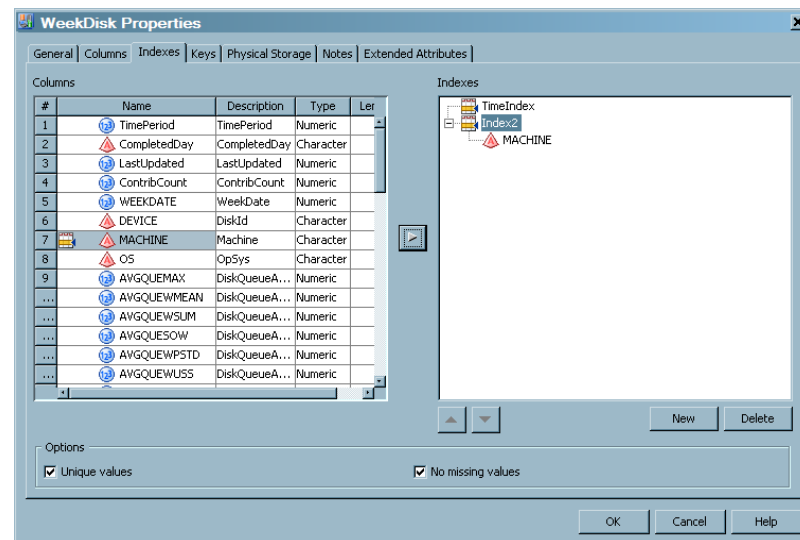
To add an index, perform the following steps:

1. Navigate to the **Index** tab of the aggregation table for which you want to create an index. For information about how to access an index, see [“How to Access an Index” on page 174](#).
2. Click **New**. An index entry (Index*n*) opens in the right panel.
Note: In the preceding display, the new index was given the default name Index2.
3. From the left panel, select the columns that you want to index. An index must have at least one column.
4. Use the right arrow to include them under the index entry in the right panel.
5. Do not specify that the values of an index are to be unique, because this might interfere with the processing of the transformation.
6. To specify that the values of an index should not include missing values, highlight the index in the right panel and click **No missing values**.
7. When you finish specifying the columns that you want to index, you might want to rename the index so that it is more informative. Otherwise, click **OK** to save your changes and return to the process flow diagram.

Note: At this time, if the name of a simple index is not the same as the column that is being indexed, a message box appears that asks if you want to change the name to conform to the name of the column. You must click **Yes** and change the name. Then you can save the index.

In the following display, the new index, called Index2, is specified for the column Machine. It contains only unique values and does not contain missing values.

Display 8.23 Newly Created Index Named Index2



Delete a Column from an Index

To delete a column from an index, perform the following steps:

1. Navigate to the **Index** tab of the aggregation table that you want to work with. For information about how to access an index, see [“How to Access an Index” on page 174](#).
2. From the right panel, expand the indexes that are displayed. Select the column that you want to delete.
3. Click **Delete**.
4. Click **OK** to save your changes and return to the process flow diagram.

Delete an Index

To delete an index, perform the following steps:

1. Navigate to the **Index** tab of the aggregation table whose index you want to delete. For information about how to access an index, see [“How to Access an Index” on page 174](#).
2. Select the index that you want to delete.
3. Click **Delete**.
4. Click **OK** to save your changes and return to the process flow diagram.

Rename an Index

To rename the index, perform the following steps:

1. Navigate to the **Index** tab of the aggregation table whose index you want to rename. For information about how to access an index, see [“How to Access an Index” on page 174](#).
2. Right-click the index entry.
3. From the drop-down list, select **Rename**.
4. Type the new name of the index. If the index contains only one column, it is a simple index. For simple indexes, the name of the index must be the same as the column that is being indexed.

Note: From this drop-down list, you can also add and delete indexes.

Press the ENTER key.

5. Click **OK** to save your changes and return to the process flow diagram. At this time, if the name of a simple index is not the same as the column that is being indexed, a message box appears that asks if you want to change the name to conform to the name of the column. You must click **Yes** and change the name. Then you can save the index.

Chapter 9

Adapter Setup Wizard

About the Adapter Setup Wizard	177
What is the Adapter Setup Wizard?	177
Advantages of the Adapter Setup Wizard	178
Naming Conventions That the Adapter Setup Wizard Uses	178
Using the Adapter Setup Wizard	179
Accessing and Navigating the Adapter Setup Wizard	179
Steps to Complete the Adapter Setup Wizard	179
Creating Jobs Using the Adapter Setup Wizard	180
What Does the Adapter Setup Wizard Create?	191
What are the Types of Jobs That the Adapter Setup Wizard Creates?	191
Where Does the Adapter Setup Wizard Store the Objects That It Generates?	193
Staging Jobs and Objects Created by the Adapter Setup Wizard	195
Aggregation Jobs and Objects Created by the Adapter Setup Wizard	196
Information Map Jobs and Objects Created by the Adapter Setup Wizard	199
Reporting Jobs and Objects Created by the Adapter Setup Wizard	199
Libraries Created by the Adapter Setup Wizard	202
Log File Created by the Adapter Setup Wizard	203
Domain Categories	204
About Domain Categories	204
Deploy and Run Jobs That the Adapter Setup Wizard Creates	205

About the Adapter Setup Wizard

What is the Adapter Setup Wizard?

The SAS IT Resource Management enhancement of SAS Data Integration Studio includes the Adapter Setup wizard. The Adapter Setup wizard enables you to define and create staging jobs, aggregation jobs, information map jobs, and reporting jobs that process the raw data for an IT resource. These jobs include predefined configurations that are based on industry intelligence that represents a best estimate about the types of aggregations and reports that most sites need for analysis.

Note: The Adapter Setup wizard produces staging, Aggregation, Information Map, and Gallery transformations that are good choices for many users. However, your site might have site-specific requirements. Therefore, you should examine the results of the Adapter Setup wizard and possibly add or delete elements of the jobs,

transformations, or tables in order to accommodate the needs of your site or organization.

After the wizard creates the appropriate jobs, you can schedule and run each as needed. The jobs process the raw data in a manner that is specified by or based on a selected adapter and domain category.

For more information about using the Adapter Setup wizard, see [“Accessing and Navigating the Adapter Setup Wizard” on page 179](#).

Advantages of the Adapter Setup Wizard

In addition to defining and programmatically creating jobs, staged tables, summarized aggregations, information maps, and reporting jobs for a given adapter, the Adapter Setup wizard also applies intelligence to these objects. Here are some of the benefits that the Adapter Setup wizard provides:

- programmatic creation of staging, aggregation, and information map jobs that include transformations and output tables that are used for staging and aggregating data with the recommended summarizations, statistics, and rankings for reporting
- programmatic creation of reporting jobs that include Gallery transformations that use predetermined information maps and report definitions to build relevant reports based on a given adapter and domain category
- options to specify the scope and size of jobs that are created for a given adapter as determined by output level (staging job only, minimal, typical, and full) or other criteria (domain category, time period, key metrics, and so on)
- translation of metrics in the raw data source to standard units of measure
- generation of a default set of computed columns that are appropriate for the adapter data source
- the ability to choose appropriate statistics for aggregations and weighting them appropriately
- a default list of metrics that will be ranked
- an appropriate set of aggregations based on the staged data for the adapter data source
- a default set of filters that might be useful for limiting the output size of aggregations
- categorization of data in summarized aggregations by appropriate specification of class variables
- creation of an information map for each aggregation table
- report-ready data

With the Adapter Setup wizard, you can create comprehensive jobs that meet your reporting needs.

Naming Conventions That the Adapter Setup Wizard Uses

The Adapter Setup wizard uses specific naming conventions when naming and saving the various folders, jobs, transformations, and tables that it creates. These standards simplify the management and maintenance of the various objects that are interrelated. For best results, if you choose to edit the IT data mart folders, subfolders, jobs, transformations, or tables that the Adapter Setup wizard creates, then use the same naming conventions that the Adapter Setup wizard uses for consistency and efficiency.

This chapter provides several examples of the naming conventions that the Adapter Setup wizard uses. For a comprehensive list of the naming conventions and location standards that SAS IT Resource Management uses, see [“About Naming Standards for SAS IT Resource Management Objects” on page 503](#).

Using the Adapter Setup Wizard

Accessing and Navigating the Adapter Setup Wizard

To access the Adapter Setup wizard, open SAS IT Resource Management and select **New** ⇒ **IT Resource Management** ⇒ **Adapter Setup**. To navigate through the pages of the wizard, use the following wizard buttons:

Help

displays a Help topic for the current window.

Cancel

abandons changes that were made since the last save and closes the current window.

Back

displays the previous window in the wizard.

Next

displays the next window in the wizard.

Finish

saves changes that were made since the last save and closes the current window.

Steps to Complete the Adapter Setup Wizard

The Adapter Setup wizard guides you through several steps that help define the ETL jobs, staged tables, summarized aggregation tables, information maps, and Gallery transformations for the selected adapter. The steps vary based on the requirements of the adapter that you select.

Here are all of the possible steps or pages that the Adapter Setup wizard can present for a given adapter:

- Specify an adapter for your input data source.
- Specify the configuration level that controls the scope and amount of output jobs.
- Specify the IT data mart.
- Specify staging parameters such as the location of the raw data and how to handle duplicate data.
- Select domain categories for the selected adapter.
- Select the time periods for which you want to create aggregation jobs and report jobs.
- Specify the location of any user-written formulas.
- Specify the SAS Content Server for report packages.
- View a summary of the specifications that you have chosen and the ETL objects that the Adapter Setup wizard creates when the resulting jobs are executed.


Creating Jobs Using the Adapter Setup Wizard

The Adapter Setup wizard enables you to define and create the ETL jobs that are required to process your IT resource data. Because the Adapter Setup wizard can create jobs without deploying them, you do not need to run the Adapter Setup wizard frequently for a given adapter. For best results, run the Adapter Setup wizard initially for a given adapter and do not run it again for the same adapter unless you have significant changes in how you would like to process input data. Running the Adapter Setup wizard additional times does not update the jobs, transformations, or tables from the previous runs. Instead, it creates new jobs, transformations, and tables.

To specify all of the necessary parameters that the Adapter Setup wizard requires to create the jobs for an adapter, perform the following steps:

1. From the SAS IT Resource Management menu, select **New** ⇒ **IT Resource Management** ⇒ **Adapter Setup** to open the Adapter Setup wizard.
2. Specify the adapter for which you want to create jobs.

Display 9.1 Specify the Adapter



- a. Click the adapter that you want to use from the list of the adapters provided. You can move the cursor over an adapter to view a tooltip that provides the full name of the adapter, any former names of the adapter, and information about the types of data that the adapter provides. The remaining pages and parameters that are presented in the Adapter Setup wizard are based on the specific adapter that you select on this page.

Note: The AS/400 and IMS adapters are not available for selection on this page. The staging code that is required for these data sources is not automated by the Adapter Setup wizard. However, SAS IT Resource Management supplies

staging transformations for these adapters. You can use these staging transformations to create staging jobs manually for these data sources. For more information about how to create a staging job manually using a staging transformation, see [“Add a Staging Transformation to a Staging Job” on page 109](#).

For more information about the specific adapters that SAS IT Resource Management supports, see [“Supported Adapters” on page 68](#).

b. Click **Next**.

3. Specify the configuration level that represents the scope of output jobs that you want the Adapter Setup wizard to create.

Note: If you are using an adapter that does not include aggregations that are categorized with configuration levels, then the level is automatically set to create a staging job only and this page does not display.

Display 9.2 Specify the Level of Output

The screenshot shows a window titled "Adapter Setup Wizard" with a subtitle "Specify the level of output". Below the subtitle is the instruction "Select the configuration level for the required output jobs." There are four radio button options: "Staging Job Only", "Minimal", "Typical", and "Full". The "Typical" option is selected. Below each option is a descriptive paragraph. At the bottom of the window are five buttons: "< Back", "Next >", "Finish", "Cancel", and "Help".

Adapter Setup Wizard

Specify the level of output
Select the configuration level for the required output jobs.

☐ Staging Job Only

☐ Minimal
Generates the staging job and a small number of jobs that include metrics, aggregations, and report definitions that satisfy a minimal set of reports for an adapter.

☒ Typical
Generates the staging job and a set of jobs that include metrics, aggregations, and report definitions that satisfy a typical set of reports for an adapter. (Includes the output that is part of the Minimal setting.)

☐ Full
Generates the staging job and a comprehensive set of jobs that include all recommended metrics, aggregations, and report definitions for an adapter. (Includes the output that is part of the Typical setting.)

< Back Next > Finish Cancel Help

Configuration levels enable you to specify the general scope or amount of aggregation tables and jobs that you want the Adapter Setup wizard to create. These levels provide various presentations of granularity in the data based on the aggregations and metrics that are available for an adapter. The jobs that are created for each level include the metrics that are used in the reports for that level. These jobs might also include other metrics that have been predetermined to be useful for reporting on the given level.

Configuration levels are unique for each adapter. However, in some cases, the levels can be identical for a given adapter. For more information about the metrics and aggregations that are supported for each adapter, see the SAS IT Resource Management metrics documentation. To locate the SAS IT Resource Management

documentation, use the Products Index at <http://support.sas.com/documentation/index.html>.

To specify the level of output, complete the following steps:

- a. Select one of the following configuration levels:

Staging Job Only

creates a staging job only for the adapter. If you select this option, no aggregation jobs, information map jobs, or reporting jobs are created.

Minimal

creates a staging job and additional jobs that include metrics, aggregations, and report definitions for a small number of reports. This level is for reporting on a small scale. You might use this level to create jobs that are related to low-priority, non-critical machines.

Typical

creates a staging job and a robust set of jobs that include metrics, aggregations, and report definitions for a typical set of reports. This level includes all of the jobs that are part of the Minimal level. In addition, it includes metrics, aggregations, and jobs that IT Resource Management has identified as being of interest for the selected adapter. You might use this level to create jobs that are related to mission critical resources.

This configuration level is the default option.

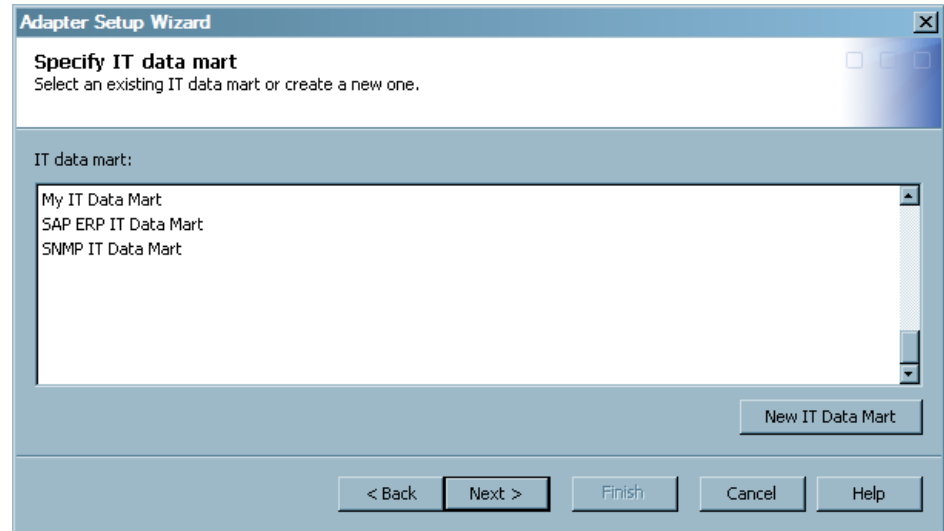
Full

creates a staging job and a comprehensive set of jobs that include most metrics, aggregations, and report definitions that SAS IT Resource Management supports. This level includes all of the jobs that are part of the Typical level. In addition, it includes many of the other metrics that are recommended for an adapter. You might use this level when you want the full benefits of the domain intelligence that is available with the IT Resource Management solution as well as additional information from the full set of metrics that are available from the adapter data sources.

Note: Increasing the configuration level results in more complex transformations and jobs. If you select **Full** when you need only the **Typical** level, then the extra complexity causes the Adapter Setup wizard to be slower to complete than is necessary for your needs. Also, the nightly ETL jobs will take longer to process.

- b. Click **Next**.

4. Specify the IT data mart where you want to save the ETL jobs that the Adapter Setup wizard will create, or create a new IT data mart.

Display 9.3 Specifying the IT Data Mart

- a. Select an IT data mart from the list of available IT data marts. If you do not want to use one of the data marts in the list, then click **New IT Data Mart**. This action launches a wizard that guides you through the process of creating and defining a new IT data mart. When you are finished, the wizard returns you to the Adapter Setup wizard.
 - b. Click **Next**.
5. Specify staging parameters for the selected adapter. Staging parameters vary by adapter as they specify the configurations that are required to extract, transform, and load specific raw data into staged tables. The following image is an example of the staging parameters for the DT Perf Sentry adapter:

Display 9.4 Specifying Staging Parameters for the Adapter

- a. Specify the staging parameters in order to govern how you want the staging code to execute for the adapter. The parameters that are marked with an asterisk (*) require a value. All other parameters do not require a value.

You can also click **Reset** for a given staging parameter to use the default value if there is one available for the parameter. If the default value is already selected, then the **Reset** option is not available. You can also click **Reset to defaults** to set all the staging parameters to the default values.

Note: **Reset to defaults** does not reset a parameter that is dependent on another parameter if the parent parameter does not change.

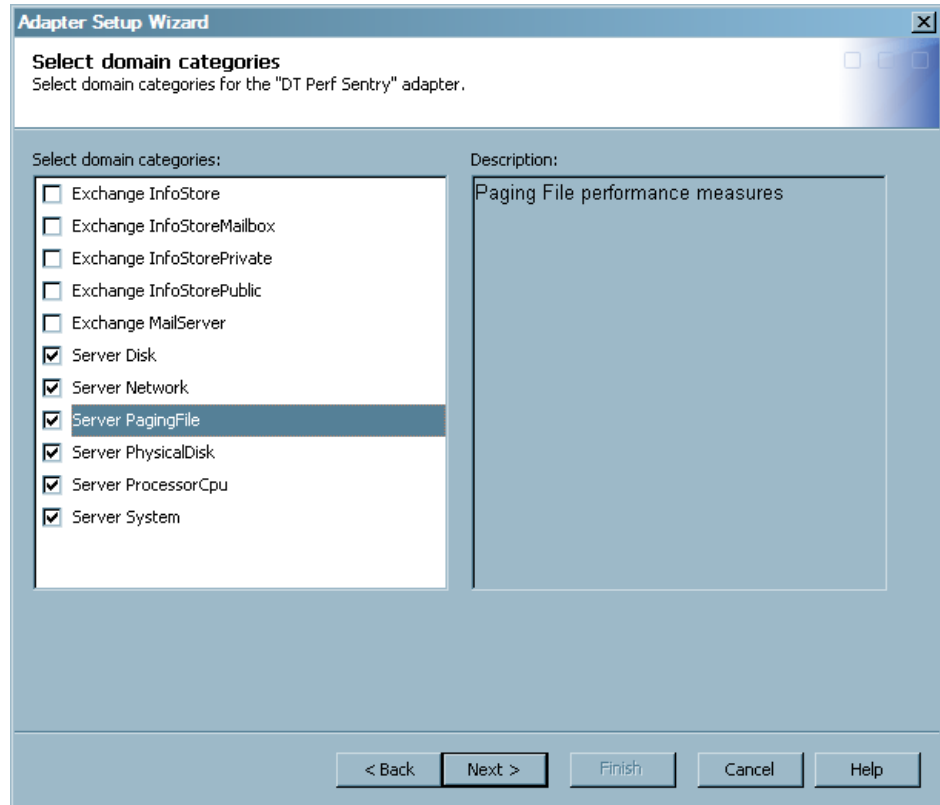
For more information about the staging parameters that are available for each adapter, see “[Staging Parameters](#)” on page 356.

- b. Click **Next**.

Note: If you selected the adapter VMware vCenter, then the Adapter Setup wizard presents two pages of staging parameters: one for VMware Data Acquisition jobs and one for VMware vCenter jobs. For more information about accessing raw data for the VMware vCenter adapter, see “[Data Sources Supported by SAS IT Resource Management Adapters](#)” on page 290.

6. Specify one or more domain categories that you want the Adapter Setup wizard to use when creating the resulting jobs.

Note: If you are using an adapter that does not have aggregation and information map jobs available from the Adapter Setup wizard, then the wizard creates only a staging job and skips this page.

Display 9.5 Select the Domain Categories for the Adapter

The preceding example shows the domain categories that are available for the DT Perf Sentry adapter. All of the domain categories that are related to Server data have been selected. Note that the description in the box on the right corresponds with the domain category that is highlighted in the **Select domain categories** box.

The Adapter Setup wizard helps you manage system performance by enabling you to select only the domain categories that are essential to the performance management and capacity planning processes for your enterprise. The Adapter Setup wizard creates a predefined set of jobs based on the domain categories that you select. The time that the Adapter Setup wizard takes to create the jobs varies based on the number of staged tables, aggregation jobs, and report jobs that are generated for the domain categories that you select.

Note: If you are using an adapter that does not include domain categories, then this page does not display.

For more information about domain categories, see [“Domain Categories” on page 204](#).

- a. Click one or more domain categories to enable or disable the domain category for staging and aggregation. At least one domain category must be selected before you can continue to the next page of the wizard.
 - b. Click **Next**.
7. Select the time periods and key metrics that you want to use for the aggregation and reporting jobs.

Note: If you are using an adapter that does not have aggregation and information map jobs available from the Adapter Setup wizard, then the wizard creates only a staging job and skips this page.

Display 9.6 Select Time Periods and Metrics

Adapter Setup Wizard

Select time periods.
Select time periods.

Specify Time Period and Shift

Select time periods:

- ☒ Day
- ☒ Week
- ☒ Month

Description:

This selection creates the aggregation jobs and the daily, weekly, and monthly report jobs that generate reports about the daily data. These jobs are based on the domain categories that you selected on the previous page. This selection includes the weekly and monthly aggregations that are needed for filtering. If you check the Shift option, additional aggregations are created and additional reports are included in the daily report jobs.

☒ Shift

Specify Key Metrics

☒ Key metrics

< Back Next > Finish Cancel Help

The preceding example shows that all time periods and key metrics have been selected. Note that the description in the box on the right corresponds with the time period that is highlighted in the **Select time periods** box.

The time periods are another way that the Adapter Setup wizard helps you manage system performance, aggregations, and jobs by enabling you to select only the time periods that are essential to the performance management and capacity planning processes for your enterprise. The Adapter Setup wizard creates only the aggregation and report jobs that reflect the time periods and key metrics that you select. In turn, you do not have to sort through extra jobs and reports that focus on time intervals that you do not need.

Note: You must select at least one time period or the **Key metrics** option to continue. (That is, you can select any combination of these four choices: **Day**, **Week**, **Month**, or **Key metrics**, but you must select at least one.)

For more information about the specific aggregations that the Adapter Setup wizard creates based on time period and shift, see [“Aggregations by Time Period” on page 196](#).

- a. Select the time periods for which you want to create aggregation and report jobs.

Day

creates aggregation jobs and the daily, weekly, and monthly report jobs that generate reports about the daily data.

Week

creates aggregation jobs and the weekly report job that generate reports about the weekly data.

Month

creates aggregation jobs and the monthly report job that generate reports about the monthly data. If you check the **Shift** option, additional aggregations are created and additional reports are included in the monthly report job.

Note: If you also select **Shift** (as described in the next step), then additional aggregations are created and additional reports are included in the report jobs.

- b. Select **Shift** if you want to generate aggregations, information maps, and reports that are based on shift data, in addition to daily, weekly, or monthly data.

Note: In some cases, **Shift** might be selected by default because of the adapter that you selected earlier in the wizard.

- c. Select **Key metrics** if you want to generate aggregation tables, information maps, and reports that include only the key metrics that focus on the most critical data for a performance area or domain category.

Note: If the domain categories that you selected do not include key metrics, then this option is not available.

- d. Click **Next**.

8. Specify the location of any user-written formulas to use in addition to the formulas that SAS IT Resource Management supplies.

Display 9.7 Specify User-Written Formulas

The staged tables and aggregation tables that the Adapter Setup wizard creates include computed columns that reference various formulas. When the tables with computed columns are instantiated, SAS IT Resource Management looks for the corresponding formulas in the list of supplied formulas.

If you have customized or user-written formulas that you want to use instead, then use this page of the wizard to specify the location of the user-written formulas that you want to use. The Adapter Setup wizard then looks in the folder that you specified for any formulas that have the same name as those in the tables that it creates. If a formula with the same name is in the folder, then the wizard uses that formula. If the formula is not found in your specified location, then the wizard searches for the formula in this location: **Shared Data/SAS IT Resource Management/IT Formulas**. If the formula is still not found, then the wizard searches for it in this location: **Products/SAS IT Resource Management/IT Formula**.

- a. Click **Browse** to navigate to and select the folder that includes any user-written formulas that you want to use instead of the formulas that SAS IT Resource Management supplies. This field is optional and does not allow manual entries. If

you do not have user-written formulas that you want to use and do not specify a folder in this field, then the Adapter Setup wizard defaults to the supplied formulas.

b. Click **Next**.

9. Specify the location of the SAS Content Server where the report definitions are retrieved and where the reports that are generated from the resulting jobs are stored. Also specify whether you want the Adapter Setup wizard to automatically deploy the jobs that it creates.

This page does not appear if you selected the configuration level **Staging Job Only** on the second page of the wizard. In this case, the Adapter Setup wizard creates only one staging job that you must deploy manually.

Display 9.8 Specify SAS Content Server and Job Deployment Page

Note: If you are using domain categories that do not result in reporting jobs, then the SAS Content Server options are not available on this page. For more information, see “Jobs That the Adapter Setup Wizard Creates for Domain Categories” on page 381.

The report definitions that the Adapter Setup wizard uses to create the report jobs are stored on the SAS Metadata Server. In addition, the reports that are generated from the report jobs are stored on the SAS Content Server.

- a. If you are using domain categories that result in reporting jobs, specify the following parameters for the SAS Content Server:

SAS Content Server

specifies the name of the server that is used as the SAS Content Server for storing report definitions and report job outputs.

Repository folder path

specifies the path within the SAS Content Server where the report definitions and reports are stored. The default location is `/SASContentServer/repository/default/sasdav/ITRM`.

URL file name parameters

URL file name parameters specifies the parameters that pertain to the URL filename, such as Proxy, user ID, or password. For example, the contents of this field could be `user='userID', pass='password'`, and so on. The contents of this field are used to communicate with Gallery Manager. In most configurations, these parameters are not needed.

A value is optional in this field. The default value for this field is blank.

For information about the contents that can be entered in this field, see the “FILENAME Statement, URL Access Method” topic in the *SAS(R) 9.3 Statements: Reference*.

Note: The contents of the field are not verified.

- b. Specify the following parameters to indicate whether you want the Adapter Setup wizard to deploy the jobs that it creates:

Deploy jobs automatically


specifies if you want the Adapter Setup wizard to deploy all of the jobs that it creates. Otherwise, the Adapter Setup wizard only creates the jobs and you must deploy them manually.

Deployment directory

specifies the location where the jobs are deployed if you select to deploy the jobs automatically.

Note: If you are working with z/OS, then the location for deployment must be a directory in the UFS file system. The Adapter Setup wizard does not support the automatic deployment of jobs to a partitioned data set (PDS) directory.

- c. Click **Next**.
10. View a summary of the metadata that you specified for the adapter, staging, aggregations, information maps, reports, and jobs that you specified with the Adapter Setup wizard.

Display 9.9 View Specifications

The summary consolidates all of the specifications that you provided on the previous pages of the wizard.

- a. Review the summary and ensure that it reflects the specifications and results that you expect. If you want to change a value that you specified on a previous page, click **Back** to navigate to the appropriate page and change the parameter value as needed.

The last item on the Summary page provides the location where a log file for the current Adapter Setup wizard instance is stored. This text file is automatically stored on your machine after each execution of the Adapter Setup wizard. The log includes all of the information that is detailed on the Summary page of the Adapter Setup wizard as well as the dates and times at which the Adapter Setup wizard was run. For more information about the log file, see [“Log File Created by the Adapter Setup Wizard” on page 203](#).

- b. Click **Finish** when complete. The wizard creates the jobs and saves them in an **<adapter><number>** folder within the IT data mart that you selected. The time that the Adapter Setup wizard takes to create the resulting jobs varies based on the adapter and the parameters that you specify. A progress bar shows the status when the wizard creates the appropriate jobs. If the jobs cannot be created, an error message is displayed.

Note: The first time you edit a job that the Adapter Setup wizard created, the job is marked as needing to be saved (that is, an asterisk is next to the name in the title bar of the Job Editor window). This is because the Job Editor is storing the initial editor settings for how the job is displayed in the window. The job was saved when the Adapter Setup wizard executed.

If you do not select to have the Adapter Setup wizard automatically deploy the jobs that it creates, then you must deploy the resulting jobs if you want to run them in

batch mode. You can right-click a job to open a menu that provides various options such as deployment and scheduling tasks. For information about running the jobs that the Adapter Setup wizard creates, see [“Running Jobs” on page 273](#). For more information about deploying a job, see [“Deploy a Job for Scheduling” on page 274](#) or SAS Data Integration Studio Help.

What Does the Adapter Setup Wizard Create?

What are the Types of Jobs That the Adapter Setup Wizard Creates?

Based on your specifications, the Adapter Setup wizard creates a set of ETL jobs and all the metadata that is necessary for those jobs to execute. These jobs and the associated metadata are stored in the <adapter><number> folder and other subfolders of the IT data mart that you are working with.

The Adapter Setup wizard creates the following types of jobs:

- one staging job that contains only the part of the ETL flow that relates to the staging transformation and the target staged tables.

Note: The Adapter Setup wizard creates two staging jobs for the VMware vCenter adapter: VMware Data Acquisition and VMware vCenter.

- zero or more aggregation jobs that each contain one Aggregation transformation with target summary tables. The number of aggregation jobs that are created varies based on the domain categories (and their respective subcategories), time periods, and reporting selections that you specified. For more information about the aggregations that the Adapter Setup wizard creates and how you can use these aggregations to analyze and report on your IT data, see [“Aggregation Jobs and Objects Created by the Adapter Setup Wizard” on page 196](#).
- zero or more information map jobs that each contain one Information Map transformation per aggregation table. The number of information map jobs that are created varies based on the domain categories (and their respective subcategories), time periods, and reporting selections that you specified.

Note: In SAS IT Resource Management 3.1 and SAS IT Resource Management 3.1.1, the Adapter Setup wizard created jobs that created or updated aggregation tables and also generated information maps based on those tables. If the job that modified aggregation tables was run, the corresponding Information Map transformations were part of that job. Therefore, the information maps were automatically regenerated. Starting with SAS IT Resource Management 3.2, jobs that create or update aggregation tables do not include transformations that generate information maps. Therefore, if any changes are made to the tables that are input to those Information Map transformations, that alter the structure of the table (for example, by adding or deleting columns), then the user must rerun the corresponding information map jobs so that the information maps are regenerated.

For more information about the information map jobs that the Adapter Setup wizard creates and how you can use them to report on your IT data, see [“Information Map Jobs and Objects Created by the Adapter Setup Wizard” on page 199](#).

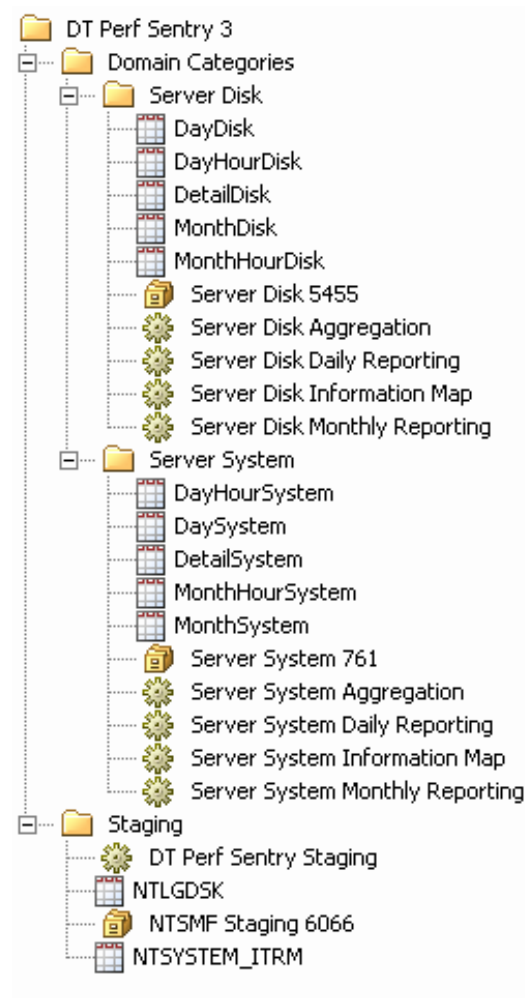
- zero or more reporting jobs that each contain one or more Gallery transformations per information map. The number of reporting jobs that are created varies based on

the domain categories (and their respective subcategories), time periods, and reporting selections that you made when specifying the adapter.

If you do not select to have the Adapter Setup wizard automatically deploy the jobs that it creates, then you must deploy the resulting jobs if you want to run them in batch mode. For more information about deploying and scheduling the jobs that the Adapter Setup wizard creates, see [“Deploy and Run Jobs That the Adapter Setup Wizard Creates” on page 205](#).

Here is an example of the subfolders and jobs that the Adapter Setup wizard created for a third instance of the DT Perf Sentry adapter with the domain categories of Server Disk and Server System selected.


Display 9.10 Subfolders and Jobs Created by the Adapter Setup Wizard



The domain categories for some adapters include a very large amount of data and associated tables. In these cases, the domain categories are divided into domain subcategories. A separate aggregation table, information map job, aggregation job, and reporting job is created for each domain subcategory based on the time periods and other reporting selections that you made when specifying the adapter. For example, the IBM SMF adapter includes a domain category called DB2 Global. The DB2 Global domain category is divided into four domain subcategories (AttachConName, ConType, JobAttach, Plan, and Subsystem).

Here is an example of the objects that the Adapter Setup wizard created for a second instance of the IBM SMF adapter with the DB2 Global domain category and the Day time period selected:

Display 9.11 Subfolders and Jobs Created for Domain Subcategories



Where Does the Adapter Setup Wizard Store the Objects That It Generates?

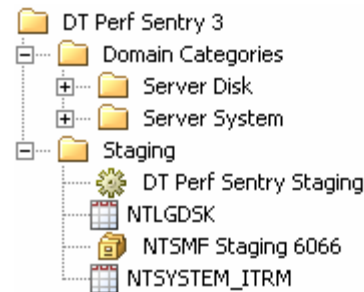
The Adapter Setup wizard stores the objects that it generates in predetermined locations within the IT data mart that you select. These objects can be stored in any location to which you have Write access. However, the Adapter Setup wizard uses the following folder organization to simplify the management and maintenance of the various objects that are interrelated. For best results, if you choose to create or edit IT data mart folders, subfolders, jobs, transformations, or tables, then use the same naming conventions and folder structure that the Adapter Setup wizard uses for consistency and efficiency.

The wizard creates an **<adapter><number>** folder in the IT data mart that you are working with. Within this folder, the wizard creates the following two subfolders for the objects that it creates:

- The **Staging** subfolder, which contains the staging job, staged tables, and the library for those tables.

The following display shows the metadata that is created in the **Staging** subfolder for a third instance of the DT Perf Sentry adapter within the IT data mart. The **Staging** subfolder contains the DT Perf Sentry Staging job, the NTLGDSK and NTSYSTEM_ITRM staged tables, and the library (NTSMF Staging 6066) for those tables.

Display 9.12 Default Location of Objects Created by the Adapter Setup Wizard for Staging the Raw Data



- The **Domain Categories** subfolder contains additional subfolders that are named by domain category. Each of these subfolders for domain categories contain the libraries, aggregation tables, and jobs (aggregation, information map, and reporting) that are associated with the respective domain category.

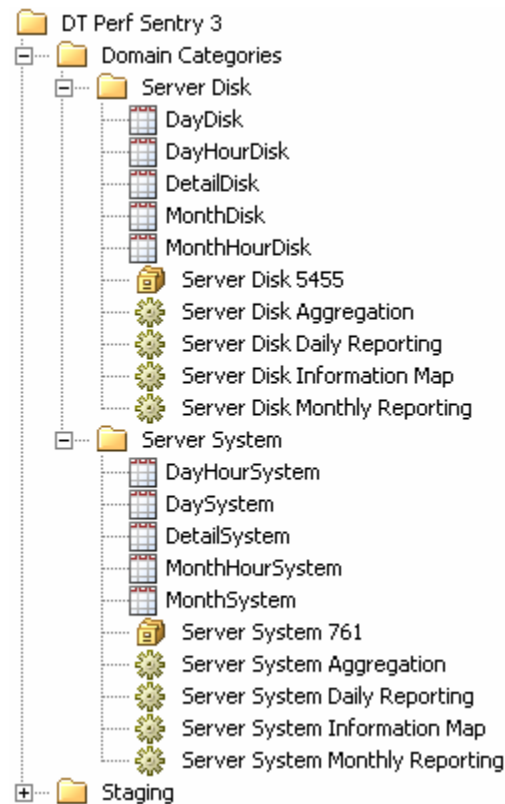
For example, the following display shows the metadata that is created in the **Domain Categories** subfolder for the following scenario:

- DT Perf Sentry adapter is selected, and it is the third time the wizard is run for this adapter in the IT data mart.
- Two domain categories (Server Disk and Server System) are selected.
- Two time periods (Day and Month) are selected.

In this example, the **Domain Categories** subfolder contains two subfolders of domain objects: Server Disk and Server System. Each domain subfolder contains five

aggregated tables, the library associated with those tables, the aggregation job, the information map job, and the two reporting jobs.

Display 9.13 Default Location of Objects Created by the Adapter Setup Wizard for Aggregating, Creating Information Maps, and Reporting on the Staged Data



Staging Jobs and Objects Created by the Adapter Setup Wizard

For most adapters, the Adapter Setup wizard creates one staging library and one staging job for each execution of the wizard. The staging library is named **<adapter> Staging <unique number>**. The staging job is named **<adapter> Staging** and it contains the following objects for a given adapter:

- one staging transformation
- one or more target staged tables

Note: For the VMware vCenter adapter, the Adapter Setup wizard creates two staging jobs (VMware Data Acquisition and VMware vCenter) with one staging transformation and one staging library for each.

A staging job creates the code that populates the associated staged tables. The staged tables are created by the Adapter Setup wizard based on template tables that are supplied as part of SAS IT Resource Management. These tables are predetermined to include the data that is most relevant and appropriate for analysis and reporting based on the domain category and time periods that you selected in the wizard. A staged table in the staging job is given the same name as the template table on which it is modeled. For more information about the specific staged tables that each domain category supports, see [“Overview of the SAS IT Resource Management Data Model” on page 459](#).

Aggregation Jobs and Objects Created by the Adapter Setup Wizard

Overview

The Adapter Setup wizard creates zero or more aggregation jobs for a supported adapter. Each of these aggregation jobs contain one staged table as the source table, one Aggregation transformation, and one or more target summary aggregation tables. The aggregation tables are named using the convention **<time period>** (such as **Day**, **WeekShift**, **MonthShiftHour**, **KeyMetrics**, **Detail**, and so on) **<domain subcategory>** (such as **cache**, **memory**, **server**, and so on). For example, an aggregation table name might be **MonthShiftHourDisk**.

There are three types of aggregations that the Adapter Setup wizard might include in an aggregation job.

regular aggregations by time period

transformations that specify one or more ways to classify, summarize, and age an input table based on a specified time period. For more information about aggregations that are created by time period, see [“Aggregations by Time Period” on page 196](#).

key metrics aggregations

aggregations that include only key metrics and any relevant classification and ID columns that are associated with these key metrics to provide vital data for the longer term forecast and capacity planning reporting. about key metrics aggregations, see [“Key Metrics Aggregations” on page 197](#).

detail aggregations

simple aggregation tables that contain detail data as it was read by the staging transformation. For more information about the detail aggregations that the Adapter Setup wizard creates, see [“Detail Aggregated Tables” on page 198](#).

The number and type of aggregations that are provided for a supported adapter varies based on the type of data that the adapter collects and the specifications that you select in the Adapter Setup wizard.

For more information about working with aggregations in SAS IT Resource Management, see [“Working with Aggregation Transformations” on page 170](#). For more information about the aggregation metrics that are provided by SAS IT Resource Management adapters, see the documentation at <http://support.sas.com/documentation/onlinedoc/itsv/metrics.html>.

Aggregations by Time Period

The Adapter Setup wizard creates a set of aggregations based on the domain categories and time periods (**Day**, **Week**, **Month**, **Shift**) that you select. These aggregations, stored in separate folders for each domain category that you selected, classify, summarize, and age an input table based on the selected time periods. For example, the **WeekDisk** aggregation includes weekly data about server disk performance. In turn, the weekly reporting job uses the information from this aggregation table and its corresponding information map to generate a report that shows the weekly performance data for the server disk.

The following table shows the aggregations that the Adapter Setup wizard creates based on selected time periods.

Table 9.1 Aggregations Created by Time Period Selection

Selected Time Period	Generated Aggregations
Day	Detail
	Day
	DayHour
Day + Shift	Detail
	Day
	DayShift
	DayHour
	DayShiftHour
Week	Week
	WeekHour
Week + Shift	Week
	WeekShift
	WeekHour
	WeekShiftHour
Month	Month
	MonthHour
Month + Shift	Month
	MonthShift
	MonthHour
	MonthShiftHour

Key Metrics Aggregations


The Adapter Setup wizard creates key metrics aggregations if you select the **Key metrics** option in the wizard. These aggregations are named with the prefix “**KeyMetrics**” and are stored with the other aggregation tables for a given domain category.

Key metrics aggregations are tables that have only a few columns that are deemed highly important for the domain category. These aggregations serve as the source for capacity planning and forecasting. These aggregations include only key metrics and any relevant classification and ID columns that are associated with these key metrics to provide vital supporting data for the selected metrics.

Key metrics are a subset of the available metrics that are provided by adapters for a given performance area such as resource utilization. You can use key metrics aggregation tables to quickly view the most critical data for a performance area.

The following image shows a key metrics aggregation, circled in red, in a process flow diagram.

Display 9.14 Key Metrics Aggregation in a Process Flow Diagram



For more information about key metrics, see [“What Are Key Metrics?”](#) on page 417.


Detail Aggregated Tables

For some adapters, the Adapter Setup wizard creates simple aggregation tables that consist of the raw data from the staged table that is appended to the simple aggregated table. These aggregated tables are named with the prefix “**Detail**” and are stored with the other aggregation tables for a given domain category.

Detail aggregated tables provide a granularity that enables you to view and report on the raw data as it is retrieved from the source. These simple aggregated tables can support join columns, computed columns, and aging, but they do not include statistics or other intelligence from SAS IT Resource Management.

The following image shows a detail aggregation, circled in red, in a process flow diagram.

Display 9.15 Detail Aggregation in a Process Flow Diagram



Information Map Jobs and Objects Created by the Adapter Setup Wizard

The Adapter Setup wizard creates one information map job (named <domain category> Information Map) for each domain category that you select. These information map jobs include an Information Map transformation for each aggregation table that the wizard creates for the domain category. Each Information Map transformation is named the same as its corresponding aggregation table <level> (such as **Day**, **WeekShift**, **MonthShiftHour**, and so on) <domain subcategory> (such as **cache**, **memory**, **server**, and so on).

Note: When the Adapter Setup wizard creates Information Map transformations, it does not populate the **Information map path** field on the **Information Map Parameters** tab. This blank field indicates that the transformation stores the information maps in the current folder where the information map job resides. This enables you to move or rename elements of the folder structure easily without having to modify each of the Information Map transformations that are affected. If you select **Browse** to modify the **Information map path** when it is blank, then the current folder location is displayed in the field.

When the information map jobs run, the Information Map transformations create information maps for the corresponding aggregation tables and the information maps are stored in the same <domain category> folder as the information map job, aggregation tables, and aggregation job. In the case of objects created by the Adapter Setup wizard, an Information Map transformation assigns the same name to the information map that it creates. As a result, the aggregation table, Information Map transformation, and resulting information map have the same name. This helps identify and associate each object easily.

Note: In SAS IT Resource Management 3.1 and SAS IT Resource Management 3.1.1, the Adapter Setup wizard created jobs that created or updated aggregation tables and also generated information maps based on those tables. If the job that modified aggregation tables was run, the corresponding Information Map transformations were part of that job. Therefore, the information maps were automatically regenerated. Starting with SAS IT Resource Management 3.2, jobs that create or update aggregation tables do not include transformations that generate information maps. Therefore, if any changes are made to the tables that are input to those Information Map transformations, that alter the structure of the table (for example, by adding or deleting columns), then you must rerun the corresponding information map jobs so that the information maps are regenerated.

The Adapter Setup wizard provides filters for the Information Map transformations that it creates. The wizard automatically provides one filter for each ranked column and one or more filters for the ranked columns of dates. These filters are specified in the Information Map transformations and propagated to the information maps that are created when the job is run. For more information about working with Information Map transformations, information maps, and information map filters, see [“About Information Maps” on page 208](#).

Reporting Jobs and Objects Created by the Adapter Setup Wizard

About Reporting Jobs Created by the Adapter Setup Wizard

The Adapter Setup wizard programmatically creates reporting jobs for the domain categories and time periods that you might select. These reporting jobs are named using the convention <domain subcategory> <time period> (such as **Daily**,

Weekly, and **Monthly**)> **Reporting**, and they are stored in the same domain category folder as the aggregation and information map jobs.

Reporting jobs include Gallery transformations that use report definitions, information maps, and aggregation tables to generate readable reports on the IT resource data. When a reporting job executes, it creates the necessary report packages and reports based on the information maps that are in the same domain category folder. The reports are then saved as tabular or graphic files (such as HTML or .png files) in the SAS Content Server that you selected in the Adapter Setup wizard. For more information about the reporting jobs that the Adapter Setup wizard creates for specific domain categories and time periods, see [“How to Determine whether the Adapter Setup Wizard Will Create Reporting Jobs for Your Selections” on page 200](#).

Note: The metadata server and SAS Content Server must be running when reporting jobs execute or the jobs will fail to generate output.

After a reporting job is deployed, scheduled, and executed, the resulting report output is available for viewing with Gallery Manager. Gallery Manager enables you to access, filter and manage your report content from a Web browser. For more information about reporting jobs and using Gallery Manager, see the *SAS IT Resource Management: Reporting Guide*.


How to Determine whether the Adapter Setup Wizard Will Create Reporting Jobs for Your Selections

SAS IT Resource Management supplies predefined report definitions for several domain categories and time periods. These report definitions include information such as the table and level, the names of the variables, the report style, and other attributes that are used to build graphical reports from your IT resource data. The Adapter Setup wizard uses these prepackaged report definitions in the reporting jobs that it creates for the domain categories and time periods that you select.

If you select a domain category and time period for which SAS IT Resource Management supplies report definitions, then the Adapter Setup wizard creates the appropriate reporting jobs. To determine whether the Adapter Setup wizard is going to create reporting jobs for the domain category and time periods that you selected, you can view the Summary page of the wizard before you select **Finish**.

The Summary page lists the report jobs that the Adapter Setup wizard will create and the report definitions that the report jobs use, as shown in the following image:

Display 9.16 Example of a Summary Page Listing Report Jobs




Note: The Summary page also includes a note if no reporting jobs will be created.

You can also determine whether the Adapter Setup wizard is going to create reporting jobs based on your selections by viewing the list of supplied report definitions for a given domain category and time period. If SAS IT Resource Management supplies report definitions for a given adapter, domain category, and time period, then the Adapter Setup wizard creates a reporting job for those selections. If there are no supplied report definitions for a given adapter, domain category, and time period, then the Adapter Setup wizard does not create reporting jobs for those selections.

To view the supplied report definitions, perform the following steps:

1. In the Folders tree, select **Shared Data** ⇒ **SAS IT Resource Management** ⇒ **3.2 IT Report Definitions**.
2. Select an **<adapter>** folder.
3. Select a **<domain category>** folder.
4. Select **Supplied**.
5. Select a **<time period>** folder.

Display 9.17 Supplied Report Definitions for DT Perf Sentry, Server Disk, Day Data**CAUTION:**

For best results, maintain copies of any supplied report definitions and any supplied report definitions that you might have modified so that your collection of reports is preserved. The Adapter Setup wizard looks in the **Supplied** folder for the report definitions to use in the reporting jobs that it creates. If you modify or remove these supplied report definitions, then you will lose or overwrite the intelligence that SAS IT Resource Management has incorporated into these definitions to ensure proper creation of the reports. Likewise, subsequent updates to SAS IT Resource Management software will overwrite any modifications that you have made to those report definitions that were supplied by the solution.

You can also create your own report definitions, Gallery transformations, and reporting jobs. You can create reporting jobs to supplement the reporting jobs that the Adapter Setup wizard creates or to add reporting jobs to the domain categories and time periods for which the Adapter Setup wizard does not create reporting jobs. Report definitions that you create can be stored in the **User Defined** folder within the domain category folder. If you decide to create your own reporting objects and jobs, use the same naming conventions and folder structure that the Adapter Setup wizard uses for consistency and efficiency. For more information about creating your own report definitions, Gallery transformations, and reporting jobs, see the *SAS IT Resource Management: Reporting Guide*.

Libraries Created by the Adapter Setup Wizard

The Adapter Setup wizard creates libraries that are required to hold the various tables that are used by the jobs that the wizard creates. The wizard creates the necessary libraries based on the adapter and domain categories that you select. For example, only a

few adapters such as IBM SMF and ASG TMON2CIC require a spin library for spin data. The Adapter Setup wizard creates a spin library only for those adapters that require one.

Staged library and libref names must be unique within the application server. The Adapter Setup wizard creates the following libraries and stores them in the same folder as their corresponding jobs and tables in the IT data mart:

library for staged tables

holds the staged tables that the Adapter Setup wizard creates. The Adapter Setup wizard creates one staged library for each staging job that it creates with every execution of the wizard. This library is named **<adapter> Staging <unique number>**. The corresponding libref name is **STG <the same unique number as the library>**.

libraries for aggregation tables

hold the aggregation tables that the Adapter Setup wizard creates. The Adapter Setup wizard creates an aggregation library for each aggregation transformation that it creates. These libraries are named **<domain subcategory> <unique number>**. The corresponding libref name is **AGG <the same unique number as the library>**.

spin library

holds the spin metadata for adapters that use spin data. These libraries are named **<adapter name> Spin <unique number>**. The corresponding libref name is **SPIN**.

Log File Created by the Adapter Setup Wizard

The Adapter Setup wizard creates a log file after each execution of the wizard and stores the text file on your machine. The last item on the Summary page of the Adapter Setup wizard provides the location where the log file is stored. On the client, these log files are located in a directory that is determined based on the APPDATA environment variable.

- On Windows XP, this location is typically **C:\Documents and Settings\userid\Application Data\SASHome\SASITResourceManagement\3.21**.
- On Windows 7, this location is typically **C:\Users\userid\AppData\Roaming\SASHome\SASITResourceManagement\3.21**.

The name of the log file is based on the adapter that you selected and the date and time that the Adapter Setup wizard was executed (**ASWlogYYMMDDHHMMSS-<adapter name> <number>**). For example, the Adapter Setup wizard might create the following log file if executed for the DT Perf Sentry adapter on June, 15 2010 at 10:35 p.m.: **ASWlog100615103500-DT Perf Sentry 1.txt**.

The Adapter Setup wizard log includes all of the information that is detailed on the Summary page of the Adapter Setup wizard as well as the processing time for each

process and the total processing time for the wizard to execute. The following image shows a sample log file from an execution of the Adapter Setup wizard.

Display 9.18 Sample Log File from the Adapter Setup Wizard

```

Adapter setup wizard processing started at: Tue Jun 15 16:46:59 EST 2010
Summary:
=====
Adapter : DT Perf Sentry
Level of output: Typical
IT data mart: My IT Data Mart
Default root path: c:\
Application server: SASApp
Operating system: win
Host: SAMPLE.COM
Port: 1234
SYSTEMROOT: C:\Program Files\SAS\

Logging the progress:
=====
Start: Getting data in the wizard..
End: Getting data in the wizard..

Start: Creating staging job...(Datetime:Tue Jun 15 16:46:59 EST 2010)
1) staging job: DT Perf Sentry Staging
Creating folders... - DONE
Creating staging library: DT Perf Sentry Staging 5147 - DONE
Reading template table(s)...- DONE
Creating staged table(s):
1) Staged table: NTLGDSK - DONE
2) Staged table: ...

4) Reporting job: Server System Monthly Reporting
Searching report definition(s) in folder '/DT Perf Sentry/Server
System/Supplied/Month'
(Datetime:Tue Jun 15 16:49:03 EST 2010)
End: Creating reporting job(s)...(Datetime:Tue Jun 15 16:49:03 EST 2010)
Refreshing IT Data Marts view... - DONE

Adapter setup wizard processing finished at: Tue Jun 15 16:49:05 EST 2010
Total elapsed time: 125 seconds

```

Domain Categories

About Domain Categories

Domain categories are categories of template tables and aggregation tables that are organized logically to represent subsets of the IT resource measurements that are available from an adapter. This organization helps you create reports from one or more related template tables and aggregation tables of IT performance measurements.

Note: The domain categories that the Adapter Setup wizard displays are not necessarily all of the available domain categories for an adapter. You can view all of the domain categories for an adapter by opening the Folders tree view and expanding **Products** ⇒ **SAS IT Resource Management** ⇒ **Adapters** ⇒ <adapter name> ⇒ **IT Template Tables**. The template tables are grouped by domain category. All of the domain categories for an adapter are in this list.

The domain categories that are associated with an adapter vary based on the requirements and capabilities of each adapter. If SAS IT Resource Management provides

aggregation and information map jobs for a given adapter, then the Adapter Setup wizard enables you to select the domain categories that you want to use. If an adapter does not have aggregation and information map jobs available from the Adapter Setup wizard, then the wizard creates only a staging job and does not enable you to select a domain category.

This capability helps you manage system performance by enabling you to select the specific domain categories that are essential to the performance management and capacity planning processes for your enterprise. The time that the Adapter Setup wizard takes to create the resulting jobs varies based on the number of staged tables, aggregation jobs and tables, information map jobs, and reporting jobs that are generated for the domain categories that you select.

Note: The Adapter Setup wizard creates ETL jobs based on the domain categories that you select. However, for some adapters such as those based on event-type data, domain categories are divided further into domain subcategories. In these cases, the jobs that are created are based on the domain subcategories instead. This prevents jobs from being too large. When working with the ETL jobs that the Adapter Setup wizard creates, you might notice that the jobs are categorized into these subcategories.

For more information about the types of jobs that the Adapter Setup wizard creates for each adapter and domain category, see [“Jobs That the Adapter Setup Wizard Creates for Domain Categories” on page 381](#).

Deploy and Run Jobs That the Adapter Setup Wizard Creates

The Adapter Setup wizard provides a convenient way to create the jobs that stage, summarize, and report on the IT resource data that an adapter loads. However, the Adapter Setup wizard does not run the jobs that it creates. You must run the resulting jobs as needed.

Deploying a job enables you to execute the job outside of SAS Data Integration Studio or the SAS IT Resource Management client, such as when scheduling a job to run at a specified time, or when a job might need to be made available as a stored process. SAS IT Resource Management enables you to deploy as well as run jobs. For more information about deploying and running jobs, see [“Jobs That Process IT Data” on page 263](#).

Because the Adapter Setup wizard creates jobs that are separated into individual component jobs, you have the flexibility to modify, deploy, and run them separately if you have updates to only one aspect of the data extraction, transformation, loading, and reporting process.

The jobs that the Adapter Setup wizard creates depend on one another to process the raw data for reporting. For example, the aggregation jobs include aggregations that rely on the staged tables that are generated when the staging job is run. As a result, each job should be run only after the jobs that it depends on have been run. The following list shows the types of jobs that the Adapter Setup wizard creates and the dependencies of each. Each job type in the list is dependent on the job type that precedes it.

1. staging job
2. aggregation job
3. information map job

4. reporting job

Chapter 10

Information Maps

About Information Maps	208
What Is an Information Map?	208
Benefits of Using Information Maps	208
How Are Information Maps Created?	208
Where Are Information Maps Stored?	209
About Information Map Transformations	210
What Are Information Map Transformations?	210
What Are the Properties of an Information Map Transformation?	210
How to Access the Properties of an Information Map Transformation	211
User-Specified Information Map Transformations	212
Overview of Specifying an Information Map Transformation	212
Create an Information Map Transformation	212
Modify an Information Map Transformation	215
Delete an Information Map Transformation from a Job	217
Delete an Information Map	218
Update an Information Map from a Template	218
Information Map Objects That Are Generated by the Adapter Setup Wizard	219
What Information Map Objects Does the Adapter Setup Wizard Generate?	219
How the Adapter Setup Wizard Names Information Map Objects	219
Where the Adapter Setup Wizard Stores Information Map Objects	220
What Information Map Filters Does the Adapter Setup Wizard Create?	220
How the Adapter Setup Wizard Names Information Map Filters	221
Working with Information Map Filters	221
Benefits of Using Information Map Filters	221
How to Access the Filters for an Information Map	222
About the Information Map Parameters Tab	222
Add a Filter to an Information Map	224
Modify a Filter on an Information Map	225
Delete a Filter from an Information Map	225
Using SAS Information Map Studio	226
Overview of SAS Information Map Studio	226
Benefits of Using SAS Information Map Studio	227
How to Access SAS Information Map Studio	227
Access an Information Map in SAS Information Map Studio	227
Create an Information Map with SAS Information Map Studio	228
Modify, Rename, or Delete Items in an Information Map with	
SAS Information Map Studio	228
Test an Information Map with SAS Information Map Studio	229

Troubleshooting Information Maps Problems	229
Information Map Does Not Appear in the IT Data Mart Folder	229
Lengthy Processing Time	230
Error Opening or Accessing Data from Information Maps	230
Errors Running Information Map Transformations	230
Information Job Fails	231

About Information Maps

What Is an Information Map?

An information map is a collection of data items and filters that describe and present a view of physical data tables in a form that is relevant and meaningful to a business user. (An information map does not contain any data—it is a map to the physical data.) You can use information maps with query and reporting applications such as SAS Enterprise Guide and SAS Web Report Studio to build business reports for your enterprise data.

Note: SAS IT Resource Management does not enable you to manipulate information maps directly. You must use Information Map transformations to specify the information maps that will be created when the transformation's job is run. As a result, the information map properties that you can view from SAS IT Resource Management are limited and not available for direct modification. These properties include metadata about the information map such as its name, location, when it was created, who created it, when it was last modified, and who modified it. They are for information purposes only and cannot be modified. If you want to edit the parameters for an information map using SAS IT Resource Management, you must modify the Information Map transformation that creates the information map and run the corresponding job. In addition, the SAS IT Resource Management solution includes SAS Information Map Studio, which you can use to modify copies of supplied information maps.

Benefits of Using Information Maps

Information maps provide the following benefits:

- Descriptive labels can be created and used rather than variable names in reports. For example, an information map might assign a column label such as “PagingFileType” to the raw data field that is named “PAGTYPE.”
- Filters are available that can be added to any information map to make it easier to report on subsets of data. For example, you can select a filter named “Last Week” rather than constructing a WHERE clause like “WeekDateDescRank=2”. (Filters can be used for multiple reports.)
- IT intelligence can be shared with other SAS applications, which can use information maps as sources for reporting. SAS Web Report Studio requires an information map as a source for its reporting processes. SAS Enterprise Guide supports information maps for reporting. SAS Foundation supports information maps by means of the information maps LIBNAME engine.

How Are Information Maps Created?

There are two ways to create an information map:

- Execute a job that contains an Information Map transformation. You can create a job that contains an Information Map transformation in two ways:
 - For many supported adapters, the Adapter Setup wizard creates jobs that contain Information Map transformations that, when run, create information maps for those adapters. For information about how to work with the Adapter Setup wizard, see [“About the Adapter Setup Wizard” on page 177](#).

Note: In SAS IT Resource Management 3.1 and SAS IT Resource Management 3.1.1, the Adapter Setup wizard created jobs that created or updated aggregation tables and also generated information maps based on those tables. If the job that modified aggregation tables was run, the corresponding Information Map transformations were part of that job. Therefore, the information maps were automatically regenerated. Starting with SAS IT Resource Management 3.2, jobs that create or update aggregation tables do not include transformations that generate information maps. Therefore, if any changes are made to the tables that are input to those Information Map transformations, that alter the structure of the table (for example, by adding or deleting columns), then the user must rerun the corresponding information map jobs so that the information maps are regenerated.

- SAS IT Resource Management enables you to create information maps for adapters by executing Information Map transformations. Information Map transformations are specified within a job. When the job runs, the Information Map transformations create information maps that are based on specified parameters. For information about how to create an Information Map transformation, see [“Create an Information Map Transformation” on page 212](#).
- Use SAS Information Map Studio to create, modify, and test information maps for SAS IT Resource Management. For more information, see the SAS Information Map Studio **Help** and [“Using SAS Information Map Studio ” on page 226](#).



Note: If you use SAS Information Map Studio to modify an information map that was created by an Information Map transformation, the changes are not carried over to the Information Map transformation. As a result, modifications that are made to an existing information map by using SAS Information Map Studio might be overwritten and lost when the corresponding job is rerun. To avoid this situation, modify a copied version of the automatically generated information map and save that modified version under a different name so that it will not be overwritten by the extract, transform, and load (ETL) processes.

Note: Information maps must be recreated whenever the structure of the tables underlying the maps is modified.

Where Are Information Maps Stored?

The Adapter Setup wizard creates folders within the IT data mart that you selected for each specified adapter. Within the adapter folder, **Domain Category** folders are set up for each separate domain that you specified. The information maps that pertain to a particular domain and the job that will create those maps are stored within that domain folder. The information map is created and stored when the job that contains the Information Map transformation is executed.

Note: When created by the Adapter Setup wizard, an information map is given the same name and is stored in the same location as the associated aggregation table. The icons that are next to the object can be used to distinguish the aggregation table from the information map. In the list of objects that can be displayed in the IT Data Marts

tree, information maps are identified by this icon: . Aggregation tables can be identified by this icon: .

When generated from a user-specified Information Map transformation, information maps are saved in any location that you specify when you set up the Information Map transformation. The information map is created and stored when the job that contains the Information Map transformation is executed.

When generated from the SAS Information Map Studio application, information maps are stored in any location that you specify when you save the new information map.

About Information Map Transformations

What Are Information Map Transformations?

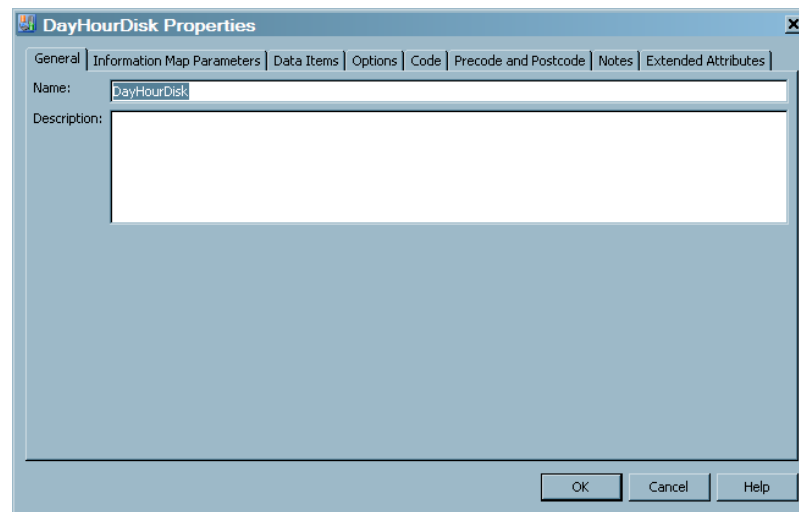
Information Map transformations are parts of a job that create information maps for source tables, which are typically aggregation tables. These transformations use PROC INFOMAPS to create information maps that describe source table items for reporting processes. Each Information Map transformation corresponds to a single source table within a job.

- When the Adapter Setup wizard runs, it automatically specifies the columns that are to become data items. The wizard also creates data items that are based on expressions, if needed.
- When you work with the Information Map transformation manually, you must select the columns that you want to include in the transformation. No columns are included automatically.

When data item specifications are added to the Information Map transformation, the name for each data item that is based on a column in the source table is derived from the label of the source column. However, these names can be changed by the user after the data item is created.

What Are the Properties of an Information Map Transformation?

The Properties dialog box of an Information Map transformation provides tabs where you can specify the parameters for the information map that it creates.

Display 10.1 Properties Dialog Box for an Information Map Transformation

As shown in the preceding display, the Properties dialog box for an Information Map transformation contains the following tabs:

- The **General** tab contains information that identifies the Information Map transformation.
- The **Information Map Parameters** tab contains the name, description, path, and filter information for the information map that this transformation will create.

Note: If the path is blank, the information map will be created in the path where the job is located.

- The **Data Items** tab contains information about the data items that are specified for this transformation, such as the source column name, the data item name that is assigned to it in the information map, the description of the data item, and the expression, if specified.

Note: A data item can be based on a source column or an expression, but not both.

- The **Options** tab contains information about additional options that can be specified for this transformation.
- The **Code** tab contains the code that is generated for this transformation.
- The **Precode and Postcode** tab contains the user-written code that can be inserted at the beginning or end of the current job or transformation. The user-written code can be stored in metadata or in a file.
- The **Notes** tab contains any notes or documents that are associated with this Information Map transformation.
- The **Extended Attributes** tab contains any property that is specific to SAS IT Resource Management and is not part of the standard metadata for an Information Map transformation.

How to Access the Properties of an Information Map Transformation

To access the properties of an Information Map transformation that are created by the Adapter Setup wizard, perform the following steps:

1. From the IT Data Marts tree, navigate to the IT data mart that contains the job where the Information Map transformation is located.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor window, and locate the Information Map transformation whose properties you want to access.
3. Right-click the Information Map transformation. From the list that is displayed, select **Properties**. The Properties dialog box appears.

Note: You have the option to store the job that contains your Information Map transformation in another location. If so, the preceding instructions are not applicable.

User-Specified Information Map Transformations

Overview of Specifying an Information Map Transformation

Using the Adapter Setup wizard is the most convenient way to create Information Map transformations and information maps for supported adapters. However, you can also specify Information Map transformations manually.

An Information Map transformation requires that you specify the following information:

- the name of the Information Map transformation
- the name of the information map that it creates
- the location where the information map is to be stored

Note: If you do not specify the location, and it is left blank, the information map is created in the same folder as the job.

- the filter specification (if filtering is needed)
- the data items that are to be included in the information map

Create an Information Map Transformation

To create an Information Map transformation, perform the following steps:

1. Navigate to the job that is to contain the new Information Map transformation. Double-click the job to open it on the **Diagram** tab of the Job Editor window.

Note: If you want to create an Information Map transformation in a new job that does not already exist, you must first create the job using the New Job wizard. For information about how to do so, see [“Jobs That Process IT Data” on page 263](#).

2. From the Transformations tree of SAS IT Resource Management, locate and expand the **ITRM Transformations** folder. Scroll down to the Information Map transformation. Then drag and drop the Information Map transformation object onto the open job on the **Diagram** tab of the Job Editor window.
3. Navigate to the IT data source for this Information Map transformation. (This source table is typically an aggregation table. However it can be any SAS table.) Drag and drop the IT data source onto the process flow diagram and connect it to the Information Map transformation object.

4. To specify the required information for the new Information Map transformation, right-click the Information Map transformation in the process flow diagram and select **Properties**. The Properties dialog box appears.
5. Click the **General** tab to enter the name and description for the Information Map transformation. The name and description that you specify here are for the Information Map transformation, not for the information map that the transformation creates.

The name of the Information Map transformation can contain no more than 60 characters.

Note: For best results, use the same naming convention that the Adapter Setup wizard uses for Information Map transformations. For information about the naming conventions that are used by the Adapter Setup wizard, see [“About Naming Standards for SAS IT Resource Management Objects” on page 503](#).

6. Click the **Information Map Parameters** tab to specify the name of the information map that is to be created and the location where the information map is to be saved.

For best results, use the following recommendations when naming the information map and determining its location:

- Use the same naming convention that the Adapter Setup wizard uses for information maps. For information about the naming conventions that are used by the Adapter Setup wizard, see [“How the Adapter Setup Wizard Names Information Map Objects” on page 219](#).
- Use the same storage path that the Adapter Setup wizard uses for information maps. To do so, do not enter a path in the **Information map path** field. (By default, the information map will be saved in the same folder as the job that contains the Information Map transformation.) For more information, see [“Where the Adapter Setup Wizard Stores Information Map Objects” on page 220](#).

However, you can change where the information map is stored by entering a path in the **Information map path** field.

Note: If **Browse** is selected and the field that contains the **Information map path** is blank, the current folder location is displayed in the **Information map path** field.

7. Use the **Information Map Filters** section of the **Information Map Parameters** tab if you want to specify filters for the information map. This section displays a grid that contains the name, expression, and description of filters that are defined for the information map that this transformation creates. (When you are creating a new Information Map transformation, the grid will be empty.) For information about how to add a filter to an Information Map transformation, see [“Add a Filter to an Information Map” on page 224](#).
8. Click the **Data Items** tab to specify the data items that you want to include in the information map that will be created by this Information Map transformation. Data items that are class columns will be available in SAS Web Report Studio for selection as filters.

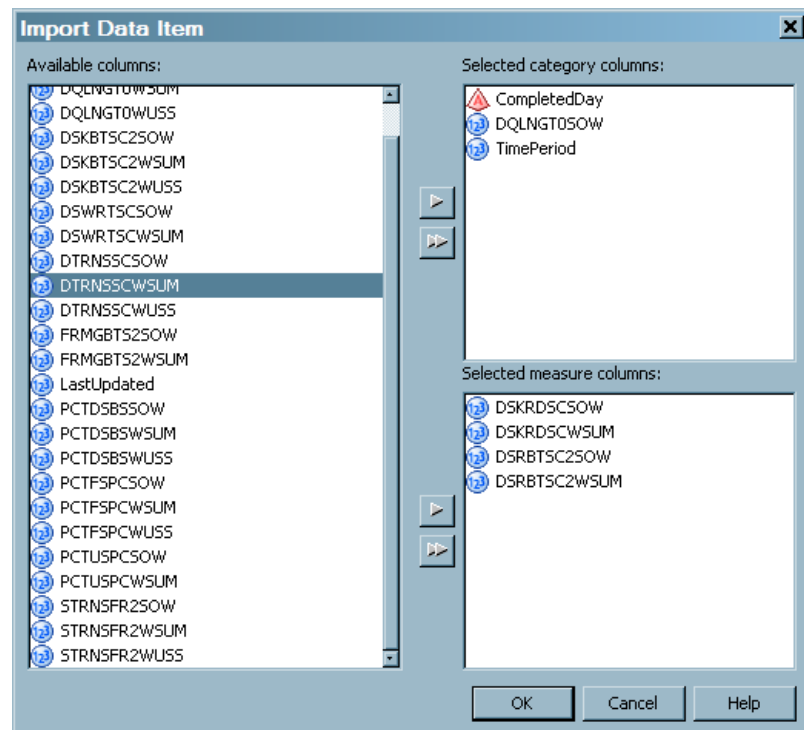
Note: In order to use or reuse the reports that are supplied by SAS IT Resource Management, specify the data items according to the standards that are in effect for the solution. For information about the standards used by SAS IT Resource Management, see [“About Naming Standards for SAS IT Resource Management Objects” on page 503](#).

The grid that appears on this tab contains columns for **Column Name**, **Data Item Name**, **Description**, **Expression**, and **Format** for each data item. When a map is being created, the grid is empty.

- Click **Import** to display the list of **Available columns**. The columns that are listed here are columns that exist in the Information Map transformation's source table. Highlight the columns that you want to include as categories in the information map and use the right arrow to transfer them to the **Selected category columns**. Highlight the columns (numeric only) that you want to include as measures in the information map and use the right arrow to transfer them to the **Selected measure columns**.

In the following display, two data items are selected as category columns, and three data items are selected as measure columns.

Display 10.2 Import Data Item Page of the Data Items Tab



Then click **OK** to display the selected category and measure columns on the **Data Items** grid where you can modify the name of the data item or its description.

- Click **New** to open the New Data Item window where you can add a new data item to the grid. Enter the **Data Item Name**, **Description**, **Classification** (Category or Measure), **Type** (Character or Numeric), **Format**, and **Expression** fields for the new data item.

Note: Do not enter double quotation marks in the **Expression** field.

In the following display, a new data item called NewMeasure is defined and ready to be added to the **Data Items** grid.

Display 10.3 New Data Item Window

The screenshot shows a 'New Data Item' dialog box. The fields are as follows:

Field	Value
Name:	NewMeasure
Description:	Description of the new measure
Classification:	Measure
Type:	Numeric
Format:	NLNUM16.2
Expression:	DURATIONSUM + 1

Click **OK** to display the new data item on the **Data Items** grid where it can be modified if needed.

Note: In the case of data items that are added from the New Data Item window, the **Column Name** of the added item is blank.

- To remove a data item from the information map, highlight the data item and click **Delete**.
9. The remaining tabs (**Options**, **Code**, **Precode and Postcode**, **Notes**, and **Extended Attributes**) are optional. For information about these tabs, see the Help for SAS IT Resource Management.
 10. Click **OK** to save the parameters for the Information Map transformation.
 11. To save the modified job, from the menu bar select **File** ⇒ **Save**.

Modify an Information Map Transformation

Editing an Information Map transformation can cause errors in reporting if the information map is used for reports. For example, do not remove or rename any data items or filters that are used in a report. If you do so, the report will fail. However, you can add anything or rename anything that is not used in a report.

To modify an Information Map transformation, perform the following steps:

1. Navigate to the job that contains the Information Map transformation that you want to modify. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Within the process flow diagram, right-click the box that represents the Information Map transformation that you want to modify.

3. From the menu list, select **Properties** to open the Information Map Properties dialog box. From this dialog box, you can access the **General**, **Information Map Parameters**, and **Data Items** tabs that serve as the primary locations for the information map parameters and metadata. These tabs enable you to view and modify the name of the transformation, the name of the generated information map, the location of the information map, the filters for the information map, and the data items that make up the information map.
4. On the **General** tab, you can modify the name and description of the Information Map transformation by changing the text in the **Name** field and the **Description** field.

Note: The name and description that you specify here are for the Information Map transformation, not for the information map that the transformation creates. For best results, use the same naming convention that the Adapter Setup wizard uses for Information Map transformations. For information about the naming conventions that are used by the Adapter Setup wizard, see [“How the Adapter Setup Wizard Names Information Map Objects”](#) on page 219.

5. On the **Information Map Parameters** tab, you can make the following changes:
 - Change the name of the information map that the transformation creates by changing the text in the **Information Map Name** field. Changing the name of an information map that was already created does not delete the map that was created with the old name.

Note: For best results, use the same naming convention that the Adapter Setup wizard uses for information maps. Furthermore, to facilitate the association of the table with the information map, the name of the information map should be the same as the name of the source table. For information about the naming conventions that are used by the Adapter Setup wizard, see [“How the Adapter Setup Wizard Names Information Map Objects”](#) on page 219.

- Change the location where the transformation saves the information map that it creates. Information maps can be stored wherever you specify. To change the location, manually enter the new path or click **Browse** to navigate to the new location. If **Browse** is selected and the **Information map path** is blank, the current folder location is displayed in the **Information map path** field.

Note: For best results, use the same storage path that the Adapter Setup wizard uses for information maps. To do so, do not enter a path in the **Information map path** field. The information map will be saved in the same folder as the information map job. However, you can directly specify where the information map is stored by entering a path in the **Information map path** field. For more information, see [“Where the Adapter Setup Wizard Stores Information Map Objects”](#) on page 220.

- Change the filters for the information map that are displayed on the filter grid.
 - For information about adding filters to information maps, see [“Add a Filter to an Information Map”](#) on page 224.
 - For information about modifying filters on information maps, see [“Modify a Filter on an Information Map”](#) on page 225.
 - For information about deleting filters from information maps, see [“Delete a Filter from an Information Map”](#) on page 225.

6. On the **Data Items** tab, you can make the following changes:

- Select additional category or measure data items from the source table to add to the information map by clicking **Import**. Then highlight the columns in the

Available columns of the source table that you want to add and use the arrow to transfer them to the **Selected category columns** or **Selected measure columns** panels. Click **OK** to include these data items in the information map.

- Define an additional data item for the information map by clicking **New**, which opens the New Data Item window. Then enter the **Data Item Name**, **Description**, **Classification** (Category or Measure), **Type** (Character or Numeric), **Format**, and **Expression** fields for the new data item.

Note: Do not enter double quotation marks in the **Expression** field.

Click **OK** to add the new data item to the information map.

- Delete a data item from the information map by highlighting it and clicking **Delete**.
7. Click **OK** to save the parameters for the Information Map transformation and the information map that it will generate. You will not be able to view the new information map that the modified transformation creates until the corresponding job has run.

Note: SAS Information Map Studio can also be used to modify the information maps that are generated by Information Map transformations in SAS IT Resource Management. However, the information maps that are generated by Information Map transformations in an information map job are regenerated each time that job is run. Therefore, any modifications that are made to an information map using SAS Information Map Studio might be overwritten and lost when the ETL job is rerun. To avoid this situation, modify a copied version of the automatically generated information map and save that modified version under a different name so that it will not be overwritten by the ETL processes.

Delete an Information Map Transformation from a Job

If you delete an Information Map transformation from a job, the information map itself is not deleted and your reports should continue to run correctly. However, if you later change a table, you must regenerate the associated information map in order to reflect the table changes in the report and thus avoid errors in the report. Do not delete an Information Map transformation unless you are certain that this action will not affect your reporting needs.

If you choose to delete an Information Map transformation from a job, perform the following steps:

1. Navigate to the job that contains the Information Map transformation that you want to delete. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
2. Right-click the Information Map transformation in the process flow diagram.
3. From the menu list, select **Delete**.
4. To save the modified job, from the menu bar select **File** ⇒ **Save**.

Note: If you delete an Information Map transformation, the information maps that it created previously are not deleted.

Delete an Information Map

Deleting an information map can create errors in reporting if your reports use the information map that you delete. Do not delete an information map unless you are certain that this action will not affect your reporting needs.

If you choose to delete an information map, perform the following steps:

1. Navigate to the location that contains the information map that you want to delete. (The default location for information maps is the folder that contains the information map job that creates the information map. However, you might have chosen to store your information map in another location.)
2. Right-click the information map. From the menu list, select **Delete**.
3. In the confirmation dialog box, click **Yes** to delete the information map. Otherwise, click **No**.

Note: If you delete an information map and not the Information Map transformation that created it, a new information map with the same name will be created the next time the corresponding job runs.

Update an Information Map from a Template

The following steps enable you to update the Information Map transformation for a table so that the information map's filters and data items conform to the supplied template. (If an information map template exists for the source table, you can use that template to automatically update the filters and data items of an information map instead of manually specifying them from the **Information Map Parameters** and **Data Items** tabs of the Properties dialog box.)

Note: You can attach the Information Map transformation to the existing aggregation job, which generated the table that will be used as the source table for the Information Map transformation. Alternatively, you can place the Information Map transformation in a new or in another existing job.

1. Open the job that is to contain the Information Map transformation.
2. In the IT Data Marts tree, locate the aggregation table that was created from a template table. Drag and drop the table onto the **Diagram** tab of the Job Editor window. (If the table was already present in this job, you can skip this step.)
3. From the Transformations tree of SAS IT Resource Management, locate and expand the **ITRM Transformations** folder. Scroll down to the Information Map transformation. Then drag and drop the Information Map transformation object onto the open job on the **Diagram** tab of the Job Editor window.
4. Connect the aggregation table to the left side of the Information Map transformation.
5. Right-click the box that represents the Information Map transformation and select **Update map definition from template**.

If an aggregation table that was created from a template is not connected to the Information Map transformation, an error message is displayed. Similarly, if an information map template does not exist for the current source table, an error message is displayed.

6. If an information map template is available for the current source table, the following message is displayed: This action will replace all the filters & data items. Do you wish to continue?
7. To update the Information Map transformation, click **Yes**.
8. When the job that contains the Information Map transformation is run, the information map will be created with the updated map attributes, including filters and data items that correspond to the template table.

Note: The Information Map transformation might have a source table that is not an aggregation table, or might have a source aggregation table that was not constructed from a supplied template. In these cases, no information map template is available. Likewise, an aggregation table that was constructed from an aggregation template, but which has been modified by the user to have a different name, might not be recognized as having an information map template.

Information Map Objects That Are Generated by the Adapter Setup Wizard

What Information Map Objects Does the Adapter Setup Wizard Generate?

The Adapter Setup wizard creates jobs that contain Information Map transformations for many of the adapters that SAS IT Resource Management supports. Using the Adapter Setup wizard is the most convenient way to create Information Map transformations and information maps for selected adapters.

When the Adapter Setup wizard is run, it creates objects that pertain to the staging and aggregating of data, and in the case of most supported adapters, the objects that pertain to the creation of information maps, such as the following:

- an information map job.
- Information Map transformations that are contained in the information map job. For most supported adapters, an Information Map transformation is created for each aggregation table.
- information maps that are generated when their corresponding Information Map transformations are run.

When generated by the Adapter Setup wizard, information map objects are defined and created according to predetermined conventions for naming, storage, and filter creation. These conventions enable you to effectively manage and locate the information map objects that the Adapter Setup wizard creates.

For more information about the Adapter Setup wizard, see [“About the Adapter Setup Wizard” on page 177](#).

How the Adapter Setup Wizard Names Information Map Objects

For consistency and efficient identification, name your Information Map transformations and information maps using the same conventions as the Adapter Setup wizard.

For most supported adapters, the Adapter Setup wizard automatically creates jobs that contain Information Map transformations. The source tables for these transformations

are aggregation tables, which are also created by the Adapter Setup wizard. The transformations are given the same names as their associated source aggregation tables—for example, `<level>` (such as week, day, month, and so on) + `<measurement category>` (such as cache, memory, server, network, and so on). The following name is an example of this naming convention:

MonthHourServicePctResponse.

The Information Map transformations also create information maps with the same name as the source aggregation tables. As a result, the aggregation tables, Information Map transformations, and resulting information maps for an information map job created by the Adapter Setup wizard all have the same name. This naming technique helps you to identify and associate each object easily.

Note: The name of the information map cannot contain any of the following characters:


* ? " ' | \ / [] \$ & > < () { } ; , ~ % ^ @ # ! In addition, the name cannot contain a newline or tab character.

Where the Adapter Setup Wizard Stores Information Map Objects

By default, the Adapter Setup wizard stores information maps in the same location as the job that contains the corresponding Information Map transformation. However, you can change where the information map is stored by revising the **Information map path** field on the **Information Map Parameters** tab of the Properties dialog box of the Information Map transformation.

For example, when the Network Information Map was run, the resulting information maps were created and stored in the default location. That is, they were stored in a subfolder named for the measurement category in the **Network** subfolder of the **Domain Categories** folder of the IT data mart. The following display shows the tree structure of these objects. The **Network** folder contains the aggregation tables, the information maps, the library, and the aggregation and information map jobs that were created by the Adapter Setup wizard.

Display 10.4 Network Domain Category



What Information Map Filters Does the Adapter Setup Wizard Create?

When the Adapter Setup wizard creates Information Map transformations, it also creates filters. The wizard automatically provides two categories of filters for columns in the aggregation tables that they are associated with.

- It provides one filter for each ranked statistic column.

- It provides one or more filters for each ranked date class column.

These filters are specified in the Information Map transformations and propagated to the information maps that are created when the job is run. The filters are then available in the information maps for analysis and reporting when you use SAS Enterprise Guide and SAS Web Report Studio.

For more information about working with information map filters, see [“Working with Information Map Filters” on page 221](#).

How the Adapter Setup Wizard Names Information Map Filters

The Adapter Setup wizard uses specific conventions for naming and describing the information map filters that it provides. These conventions use best practices for naming and describing filters in a way that is concise, meaningful, and effective for use with other SAS solutions. For best results, use the following conventions that the Adapter Setup wizard uses when naming and describing the information map filters that you create:

- Provide each filter with a unique name within an information map. No two filters in an information map can have the same name. However, filters with the same name can be in multiple information maps.
- Specify filter names that do not exceed 32 characters in length. The information map library engine that is used by SAS Enterprise Guide to access information maps truncates filter names that are longer than 32 characters.
- Include the ranked column name and the corresponding filter value in the description for the filter. The Adapter Setup wizard includes the ranked column name and the corresponding filter value in the description for the filters that it provides. For example, if the **CacheReadHitPctMinAscRank** column is filtered in ascending order from one to ten (including one and ten), then the description for that filter is **“CacheReadHitPctMinAscRank 1-10.”** For best results, use this convention to describe the information map filters that you create so that your information map filters are consistently identifiable.

Working with Information Map Filters

Benefits of Using Information Map Filters

Information maps support filters that enable you to subset data for efficient reporting. These filters make reporting easier by filtering out data that is not necessary for a given report and focusing on the data of interest. For example, you might create and use a filter named “Last Week” that specifies a WHERE clause for **WeekDateDescRank=2** in order to report on last week's data. Or, you might create a filter named “E-mail Servers” to filter the data with a WHERE clause **Domain='E-mail'** in order to report on a subset of the data that relates to an e-mail server.

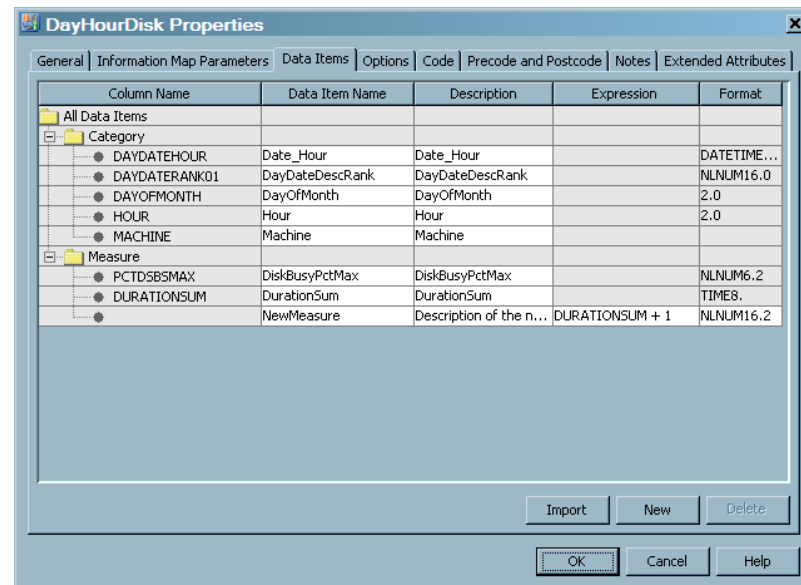
As mentioned in the preceding topic, the Adapter Setup wizard automatically provides several filters for the Information Map transformations that it creates. If you create your own filter for an information map, ensure that the column with which you want to create the filter has been defined in the source table.

Note: A report can use one or more filters and these filters can be selected at reporting time. Unlike aggregation filters, these filters do not filter out any data until they are used in a report.

How to Access the Filters for an Information Map

The filters for an information map can be accessed on the **Information Map Parameters** tab of the Properties dialog box of the Information Map transformation. To access the filters, perform the steps that are shown in “[How to Access the Properties of an Information Map Transformation](#)” on page 211. Then select the **Information Map Parameters** tab.

Display 10.5 Information Map Parameters



About the Information Map Parameters Tab

As shown in the preceding display, the following information is displayed on the **Information Map Parameters** tab.

Information map name

specifies the name of the information map that the transformation creates. The name can contain no more than 60 characters. The name of information maps cannot contain any of the following characters: * ? " ' | \ / [] \$ & > < () { } ; : ` ~ % ^ @ # ! In addition, information map names cannot contain a newline, or tab character. A value is required in this field.

Information map description

specifies information about the information map. A value is optional in this field.

Information map path

specifies where the transformation stores the information map that it generates. Click **Browse** to select a path to the information map from the Select a Location dialog box that displays the metadata directory structure.

Note: This field can be either an absolute path or blank. If the field is blank, the information map will be stored in the same folder as the job that contains the

Information Map transformation. If **Browse** is selected and the **Information map path** is blank, the current folder location is displayed in the **Information map path** field.

Information Map Filters

lists the filters for the information map in the form of a grid. Filters in an information map are criteria (rules) that subset data. You can click any of the fields of a filter (**Filter Name**, **Filter Expression**, and **Filter Description**) to change them. You can also add and delete filters for the information map.

#

The numbers under this symbol specify the ordinal position of the filter for the information map.

Filter Name

specifies the name of the filter. The name must be unique and can contain up to 60 characters.

Filter Expression

specifies the filter criteria (rules). The valid form for an expression of an information map filter is *<<name of the table.name of the column>> valid operator and value for a numeric column or a string column.*

CAUTION:

Filter expressions for information maps are case-sensitive. Filter expressions must not contain any double quotation marks. A discrepancy in case might cause errors or unexpected results when using the information map filter with other SAS solutions. Filter expressions that contain double quotation marks might cause the Information Map transformation to fail when executed.

Here are three examples of valid expressions for information map filters:

- *<<DaySystem.DayDateDescRank>> = 1*
- *<<DaySystem.DayDateDescRank>> between 1 and 7*
- *<<DaySystem.Machine>> contains 'MACHINENAME'*

Note: For more information about filter expressions, see the documentation for PROC INFOMAP in the *Base SAS Procedures Guide* at <http://support.sas.com/documentation/onlinedoc/base/index.html>.

Description

displays a description of the filter. The description can contain up to 200 characters. However, short descriptions require less system memory and are often processed more efficiently by various SAS solutions.

Note: The Adapter Setup wizard includes the ranked column name and the corresponding filter value in the description for the filters that it provides. For example, if the **CacheReadHitPctMinAscRank** column is filtered in ascending order from one to ten (including one and ten), then the description for that filter is **"CacheReadHitPctMinAscRank 1-10."** For best results, use this same convention to describe the information map filters that you create.



reorders filters by moving a selected filter up in the list.



reorders filters by moving a selected filter down in the list.

New

This button enables you to enter a new information map filter. When you click **New**, a new row is displayed in the filter list and you can enter the parameter values for the new filter.

Delete

This button enables you to delete an information map filter. Highlight the filter that you want to delete and click **Delete**.

Add a Filter to an Information Map

To add a filter to an information map, perform the following steps:

1. Navigate to the location that contains the information map where you want to add a filter.
2. In the appropriate **Domain Categories** folder of the adapter that you are working with, locate the job that contains the transformation for the information map.
3. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
4. Right-click the box that represents the Information Map transformation. From the menu list, select **Properties** to open the Properties dialog box.
5. Select the **Information Map Parameters** tab to view a grid of filters. This grid displays the name, expression, and description of filters that are already defined for the information map that this transformation creates.
6. At the bottom of the page, click **New**. This action displays a new line in the grid on which you can enter a new filter. The default name of this filter is “Untitled n ,” where n is the n th filter that exists for this information map.
7. In the **Filter Name** column, enter the name of the filter. This name must be unique within the information map.

Note: Although the **Filter Name** field can contain up to 60 characters, the information map library engine that is used by SAS Enterprise Guide to access information maps truncates filter names that are longer than 32 characters. To avoid the potential for truncation when using this information map with other SAS solutions, specify a filter name that does not exceed 32 characters in length.

8. In the **Filter Expression** column, enter the expression that specifies the filter criteria. The valid form for an expression of an information map filter is *<<name of the table.name of the column>> valid operator and value for a numeric column or a string column*. Here are three examples of valid expressions for information map filters:

- <<DaySystem.DayDateDescRank>> = 1
- <<DaySystem.DayDateDescRank>> between 1 and 7
- <<DaySystem.Machine>> contains 'MACHINENAME'

If you enter an invalid expression, the message “Error: Failed to insert filter <name-of-filter>” will display in the log when the job is run.

CAUTION:

Filter expressions for information maps are case-sensitive. Filter expressions must not contain any double quotation marks. A discrepancy in

case might cause errors or unexpected results when using the information map filter with other SAS solutions. Filter expressions that contain double quotation marks might cause the Information Map transformation to fail when executed.

Note: SAS Enterprise Guide honors both measures and categories for information map filters. In SAS Enterprise Guide, you can open the information map that is generated from this transformation and you can select the filters that you have created.

9. In the **Description** column, enter a brief description of the filter. The description can contain up to 200 characters. However, short descriptions require less system memory and are often processed more efficiently by various SAS solutions. For best results, minimize the number of characters in the description while maintaining meaning and clarity.

Note: The Adapter Setup wizard includes the ranked column name and the corresponding filter value in the description for the filters that it provides. For example, if the **CacheReadHitPctMinAscRank** column is filtered in ascending order from one to ten (including one and ten), then the description for that filter is “**CacheReadHitPctMinAscRank 1–10.**” For best results, use this same convention to describe the information map filters that you create.

10. Click **OK** to save. The new filters that you create are available in the information map only after the corresponding job runs.

Modify a Filter on an Information Map

To modify a filter on an information map, perform the following steps:

1. Navigate to the location that contains the information map where you want to modify a filter.
2. In the appropriate **Domain Categories** folder of the adapter that you are working with, locate the job that contains the transformation for the information map.
3. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
4. Right-click the box that represents the Information Map transformation.
5. From the menu list, select **Properties** to open the Properties dialog box.
6. Select the **Information Map Parameters** tab to view a grid of filters. The grid displays the name, expression, and description of filters that are defined for the information map that this transformation creates.
7. In the **Information Map Filters** panel of the tab, highlight the field of the filter that you want to change and enter the revised value. In addition, you can click **New** to add a filter to the grid or **Delete** to delete a highlighted filter from the grid.
8. Click **OK** to save the Information Map transformation with its revised filter. The filter is revised only after the corresponding job runs.

Delete a Filter from an Information Map

Deleting an information map filter can create unexpected results in reporting if your reports use the filter that was deleted. For example, if a report used an information map filter that was deleted, then the next time the report is run, it can experience errors or produce unexpected results. Do not delete an information map filter unless you are certain that this action will not negatively affect your reporting needs.

To delete a filter from an information map, perform the following steps:

1. Navigate to the location that contains the information map where you want to delete a filter.
2. In the appropriate **Domain Categories** folder of the adapter that you are working with, locate the job that contains the transformation for the information map.
3. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
4. Right-click the box that represents the Information Map transformation.
5. From the menu list, select **Properties** to open the Properties dialog box.
6. Select the **Information Map Parameters** tab to view a grid of filters. The grid displays the name, expression, and description of filters that are defined for the information map that this transformation creates.
7. Right-click the filter that you want to delete.
8. From the menu list, select **Delete**.
9. Click **OK** to save. The filter is deleted in the information map only after the corresponding job runs.

Note: Filter changes are not applied immediately. They are applied the next time the job runs.

Using SAS Information Map Studio

Overview of SAS Information Map Studio

SAS Information Map Studio enables you to create and manage information maps outside of SAS IT Resource Management and SAS Data Integration Studio. When you are working with information maps for SAS IT Resource Management, you might choose to use SAS Information Map Studio to perform the following tasks:

- Create and maintain additional information maps other than those that are created by Information Map transformations in SAS IT Resource Management.

Note: Information maps can be saved in any location that you specify. You might also choose to use the same guidelines that the Adapter Setup wizard uses to manage information maps so that all of your information maps are consistent for your IT resource data.

- View information maps that were created by Information Map transformations in SAS IT Resource Management.

Note: You should not use SAS Information Map Studio to modify information maps that were created by Information Map transformations. Changes that are made to an information map using SAS Information Map Studio do not persist in the SAS IT Resource Management Information Map transformation that created the information map. As a result, the changes that are made in SAS Information Map Studio might be overwritten the next time the information map job runs. To avoid this situation, modify a copied version of the automatically generated information map and save that modified version under a different name so that it will not be overwritten by the ETL processes.

- Test information maps that were created by Information Map transformations in SAS IT Resource Management.

Benefits of Using SAS Information Map Studio

SAS Information Map Studio provides advanced capabilities for creating and managing information maps. These capabilities can provide some advanced benefits when working with information maps for your IT resource data. Here are some of the advanced features that are available in SAS Information Map Studio that might be helpful when managing your information maps:

- joining multiple tables within an IT data mart, across multiple IT data marts, and with other tables outside of the IT data mart.

For information about the rules that govern joins, see [“Specify Join Columns” on page 151](#).

- defining prefilters in information maps.
- using only a subset of available columns to make reporting easier by providing fewer choices of appropriate columns to report on.
- creating folders within an information map to organize data items and filters.

How to Access SAS Information Map Studio

To access SAS Information Map Studio and connect to the metadata server that contains your data, perform the following steps:

1. Open SAS Information Map Studio. To do so, from your Start menu, select **Programs ⇒ SAS ⇒ SAS Information Map Studio**
2. If you are not connected to a metadata server, perform these steps:
 - a. On the SAS Information Map Studio menu bar, right-click **File**.
 - b. From the drop-down list, select **Connection Profile** where you can specify metadata server that contains the tables that you want to work with.
 - c. When the Connection Profile dialog box appears, you can create a connection profile with the Connection Profile wizard. Alternatively, you can open an existing one by selecting the profile from the drop-down list.
 - d. If you want to use this connection profile as your default profile, click the corresponding option.
 - e. Click **OK** to connect to the metadata server.


Access an Information Map in SAS Information Map Studio

You can use SAS Information Map Studio to create, modify, view properties, and test information maps.

CAUTION:

If the information map was created by a SAS IT Resource Management Information Map transformation, save it with a different name. Changes that are made to an information map using SAS Information Map Studio do not persist in the SAS IT Resource Management Information Map transformation that created the information map. As a result, the changes that are made in SAS Information Map

Studio might be overwritten the next time the information map job runs. To avoid this situation, modify a copied version of the automatically generated information map and save that modified version under a different name so that it will not be overwritten by the ETL processes.

1. After you connect to the metadata server, use the Folders tree () in the Resources panel to navigate to the location of the information map that you want to work with.
2. Highlight the information map that you want to open and click **Open**. Alternatively, you can double-click the information map to open it.
3. The Properties panel displays the properties and the values of the object that is highlighted in the Information Map Contents panel. Click **Details** to display more information about the highlighted object.
 - To view the properties of the information map, highlight the information map name in the Information Map Contents panel.
 - To view the properties of a column, highlight the column name in the Information Map Contents panel.

For more information about SAS Information Map Studio, see the SAS Information Map Studio Help and *SAS Information Map Studio: Getting Started with SAS Information Maps*.

Create an Information Map with SAS Information Map Studio

To create an information map, select **File** ⇒ **New** from the menu bar. In the Resources panel, navigate to the IT data mart that contains the table for which you want to create an information map. Double-click the table to display its items in the Information Map Contents panel.

- To add a new folder, data item, or filter to the map, in the Information Map Contents panel right-click the map and select an option from the drop-down list.

Note: Do not enter blank spaces in the **Name** field of a data item. To do so causes errors when the report is processed in Gallery Manager.

- To delete an item from the map, highlight the item and select **Delete** from the drop-down list.
- To rename an item on the map, highlight the item and select **Rename** from the drop-down list.

Save the information map by selecting **File** ⇒ **Save** from the menu bar.

Note: To be consistent with the naming conventions of the Adapter Setup wizard, save the information map under the same name as the table on which it is based. For information about the naming conventions that are used by the Adapter Setup wizard, see [“How the Adapter Setup Wizard Names Information Map Objects” on page 219](#).

Modify, Rename, or Delete Items in an Information Map with SAS Information Map Studio

To modify an existing information map, navigate to the map and double-click it. The items in the existing map are displayed in the Information Map Contents panel.

- To add a new folder, data item, or filter to the map, in the Information Map Contents panel, right-click the map and select an option from the drop-down list.

Note: Do not enter blank spaces in the **Name** field of a data item. To do so causes errors when the report is processed in Gallery Manager.

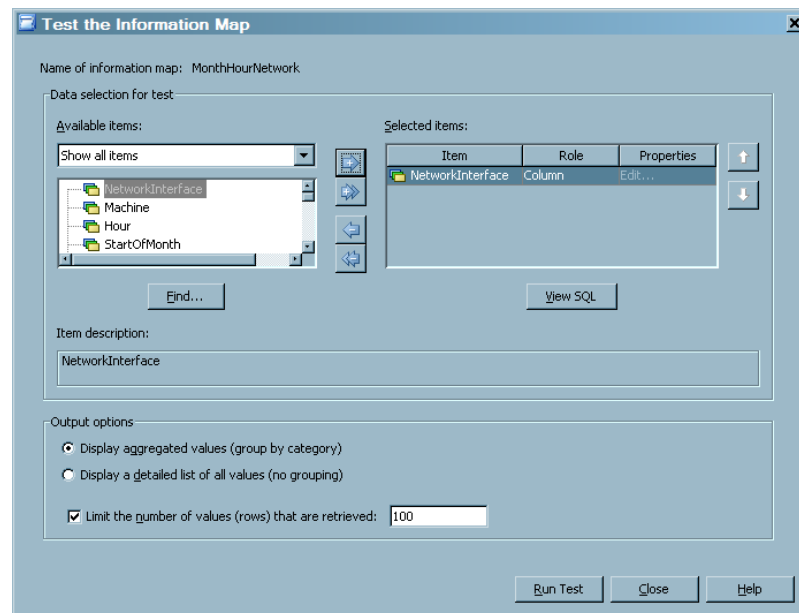
- To delete an item from the map, right-click the item and select **Delete** from the drop-down list.
- To rename an item on the map, right-click the item and select **Rename** from the drop-down list.
- When you are finished making changes, click **OK**.

Save the information map by selecting **File** ⇒ **Save** from the menu bar.

Test an Information Map with SAS Information Map Studio

To test a map, select **Tools** ⇒ **Run a Test Query**. Then, select one or more items from the **Available items** panel and add them to the **Selected items** box.

Display 10.6 Test the Information Map Window



Set the Output options to specify how you want to view the results and click **Run Test**. The results are displayed in a separate window.

Troubleshooting Information Maps Problems

Information Map Does Not Appear in the IT Data Mart Folder

Problem: After successful execution of an Information Map transformation, the resulting information map is not displayed in the IT Data Mart folder.

Solution: To view the results of a successful execution of an Information Map transformation, you must explicitly refresh the current tree or work area. To do so, right-click somewhere in the tree. From the drop-down menu, click **Refresh**.

Lengthy Processing Time

Problem: Information maps take quite a bit of time to process.

Solution: To save processing time, you might want to delete any Information Map transformations for information maps that will not be used. You might also save processing time by deleting filters that are not useful.

Error Opening or Accessing Data from Information Maps

Problem: An error occurs when opening or accessing data from information maps when using SAS Enterprise Guide, SAS Web Report Studio, or SAS Information Delivery Portal.

Solution: To determine whether you have Read access restrictions that prevent you from accessing the data, perform the following steps:

- Open SAS Information Map Studio. Navigate to the information map that experienced the error.
- From the Resources panel, select items from the information map that you want to verify and add them to the Selected Resources panel. Verify that the columns of the selected table are displayed correctly in the Information Map Contents panel.
- Select **Tools** ⇒ **Run a Test Query**.
- Click **Run Test**.

If you see any errors or warnings during the preceding steps, verify that you have Read privileges and Read metadata privileges. Otherwise, contact your data administrator.

Errors Running Information Map Transformations

Problem: An error occurs when running the Information Map job and information maps are not created.

Solution: Information Map transformations use PROC INFOMAPS, which requires specific authorization settings. Review the following list of authorization requirements and recommended configurations to ensure that your settings are configured correctly:

- Your user logon ID and password must be stored in metadata, using the SAS Management Console.

Note: At run time, PROC INFOMAPS authenticates to the metadata server using the one-time password technique documented for SAS 9.2 and later. This means that the user submitting the code (whether interactively or in batch) should have a User object defined in metadata and a Login object defined for each authentication domain, with the correct user ID and password for each. In SAS Management Console, the **User Manager** can be used to create and modify Users. For each User object, the Logins can be created or modified on the **Accounts** tab of the User dialog box. For more information, see the “Security Overview” chapter in the *SAS Intelligence Platform: Overview*.

- You must have access to the physical files and directories. (PROC INFOMAPS does not validate the table at run time.)
- You must have read authorization to the physical tables.
- You cannot have multiple SAS libraries with the same metadata name.

- Your default ACT on SAS Foundation must be set according to the configurations that are detailed in the installation guide for your operating environment.

Note: You can submit jobs if you are logged on as an administrative user. However, this is not recommended.

Information Job Fails

If you change the name of the source aggregation table or the name of a column in that table, you must manually change the associated filter expressions for the information map. You can then redeploy and rerun the job.

Chapter 11

User-Written Staging Code

User-Written Staging Transformations	233
Overview	233
What Are the Properties of the User-Written Staging Transformation?	234
General Tab	234
Staging Parameters Tab	234
Code Tab	235
Precode and Postcode Tab	235
Notes Tab	235
Extended Attributes Tab	235
Guidelines for Processing Data with User-Written Staging Transformations . . .	235
About Processing Data with User-Written Staging Transformations	235
Staging Data with User-Written Staging Transformations	236
Aggregating Data Created by User-Written Staging Transformations	247
Creating Information Maps for User-Written Adapters	249
Reporting on Data Created with User-Written Staging Transformations	252
Working with User-Written Staging Transformations	252
Create a User-Written Staging Transformation	252
Delete a User-Written Staging Transformation from a Job	258
Modify a User-Written Staging Transformation	259
Rename a User-Written Staging Transformation	259
Using the Additional Generated Code	259
Macro Variables	260
Macros	260
Tabinfo and Varinfo	261

User-Written Staging Transformations

Overview

The SAS IT Resource Management User-Written Staging transformation provides a place for you to enter your staging code for a data source. The process consists of the following steps:

1. dropping the User-Written Staging transformation onto a job on the **Diagram** tab of the Job Editor window
2. (optional) attaching a source table

3. writing the SAS program to load raw data into SAS tables
4. attaching target tables
5. specifying the options and parameters for the transformation, including the location of the staging code
6. generating the code and running the job

When the code is generated, macros and macro variables that might be useful in the user-written code are also generated. After reviewing the generated code, you can save the job and run it to create the physical tables. However before doing so, you must create the metadata for the target table by invoking the New Table wizard. (To do so, from the SAS Data Integration Studio menu, select **File** ⇒ **New** ⇒ **Table**.)

For more guidance about the processes that enable you to create user-written staging transformations, see [“Guidelines for Processing Data with User-Written Staging Transformations” on page 235](#). This topic also includes information about aggregating, creating information maps, and reporting on the data that is produced by your user-written staging code.

The remainder of this chapter consists of the four following topics:

- [“What Are the Properties of the User-Written Staging Transformation?” on page 234](#)
- [“Guidelines for Processing Data with User-Written Staging Transformations” on page 235](#)
- [“Working with User-Written Staging Transformations” on page 252](#)
- [“Using the Additional Generated Code” on page 259](#)

What Are the Properties of the User-Written Staging Transformation?

General Tab

The **General** tab contains the following information about the table. The fields on this tab can be modified.

Name

identifies the name of the user-written staging transformation that you are specifying.

Description

describes the user-written staging transformation that you are specifying.

Staging Parameters Tab

The **Staging parameters** tab contains three subgroups of information: Staging, Duplicate Checking, and User Written. In the left panel, you can select the subgroup whose properties you want to specify. For information about these properties, see [Step 8 of the Create a User-Written Staging Transformation topic on page 253](#).

Code Tab

The **Code** tab can display and manage the code that is generated for the User-Written Staging transformation. Click **Help** for more information about this tab.

Precode and Postcode Tab

The **Precode and Postcode** tab contains **Precode** and **Postcode** check boxes that enable you to specify that user-written code should be inserted at the beginning or end of the current job or transformation. Click **Help** for more information about this tab.

Notes Tab

The **Notes** tab contains areas where you can annotate the table. Click **Help** for more information about this tab.

Extended Attributes Tab

The **Extended Attributes** tab contains custom properties that are not part of the standard metadata for the transformation. All extended attributes that are added to this tab will be made into macro variables when the code for this transformation is generated. Click **Help** for more information about this tab.

Guidelines for Processing Data with User-Written Staging Transformations

About Processing Data with User-Written Staging Transformations

The user-written staging transformation enables you to create your own staging code. You can process data from any data source that SAS can read and you can generate user-defined staged tables. These user-defined staged tables can then be used as the source tables for additional transformations, including the Aggregation and Information Map transformations, in preparation for reporting on the data in those tables.

The following topics provide guidelines that describe how to work with user-written transformations:

- [“Staging Data with User-Written Staging Transformations” on page 236](#)
- [“Aggregating Data Created by User-Written Staging Transformations” on page 247](#)
- [“Creating Information Maps for User-Written Adapters” on page 249](#)
- [“Reporting on Data Created with User-Written Staging Transformations” on page 252](#)

These topics explain the process of developing a user-written adapter to be used with the SAS IT Resource Management 3.2 client. The method described in the following topics is not the only way to create an adapter. You can write an adapter in any way that you want. However, this method can help you follow the same standards that are used for the supplied SAS IT Resource Management 3.2 adapters. Following these standards keeps

your adapter consistent with other adapters. You can use some or all of the following processes as needed.

Staging Data with User-Written Staging Transformations

Overview of Staging Data with User-Written Staging Transformations

The staging step reads data from a raw data source and makes SAS staged tables. These staged tables are then used as the source to the Aggregation transformation. The User-Written Staging transformation is provided in the SAS IT Resource Management client to help in the development of the staging process. This transformation enables staged table developers to point to their own SAS code that can read the raw data and create the SAS tables. In addition, the User-Written Staging transformation supplies a set of macros and macro variables that provide information from the metadata that can be useful when creating the target tables.

The following topics discuss best practices or guidelines for staging the data with user-written staging transformations, aggregating the data processed by those transformations, creating information maps for the aggregated tables of data, and reporting on that data.

For more information about staged tables, see [“About Staging the Data” on page 106](#).

How to Stage Data with User-Written Staging Transformations

Use the following steps as a guide to staging the data with the User-Written Staging transformation.

1. Design the adapter.

To accomplish this step, you need to thoroughly understand the raw data. You also need to determine what your staged tables should look like. For example, you should determine the names of the staged tables, the column names for each table, the attributes of the columns, and so on.

2. Create an IT data mart.

On the **IT Data Marts** tab create an IT data mart for your adapter. For information, see [“Create an IT Data Mart” on page 57](#).

3. Set up the metadata for your staged tables.

Create the folder for your staged tables and libraries.

- a. Right-click the IT data mart that you just created. Select **New** ⇒ **New Folder**. Create the folder. The name of the folder should be the name of your adapter.
- b. Right-click the new folder that you just created and select **New** ⇒ **New Folder**. Create a folder called **Staging**.

Create a library and its metadata for your staged tables.

- c. Right-click the **Staging** folder and select **New** ⇒ **Library**. The New Library Wizard opens.
- d. Select **SAS Base Library** as the type of library that you want to define. On the next page of the wizard, specify a name such as Staging for the new library.
- e. On the next page of the wizard, choose SASITRM.
- f. On the next page of the wizard, specify the libref that is associated with the staged table library.

For best results, specify the name of the libref as *stgxxxx*, where *xxxx* is a random number to ensure that the libref is not used elsewhere.

In the Path Specification box, click **New** to add a new path.

TIP Use the root path that you specified for the IT data mart that you created and add an identifier to the end of that path. The identifier could be the folder, directory, or name of the libref, according to the standards of your operating environment. For best results, use the same name as the libref you just specified. For example, on Windows, if your root path is **C:\DataMarts\MyAdapter**, and your libref is *stg1357*, then specify this library path as **C:\DataMarts\MyAdapter\stg1357**.

- g. Continue to follow the prompts of the wizard to create the appropriate metadata for the library.
4. On your operating system, create the physical folder for the library. (The previous step created only the metadata for the library.)
5. Write the SAS code to read the raw data and create SAS tables as output. The complexity of this step varies depending on the raw data that you are reading and the tables that you are creating.

TIP This code can be any SAS program that creates one or more permanent SAS data tables as output. These tables should be put in the location that you created in step 4. It will simplify your work later if you use the same libref that you specified in step 3.

For best results, consider adding the following columns to your staged tables.

- **DATETIME**

In the supplied ITRM adapters, the staged tables are date and time based. All the tables contain a column called **DATETIME** to hold this information. Using this column name enables you to use the supplied formulas and the duplicate and future checking processes without modification. Its attributes are as follows:

```
attrib DATETIME length = 8
              format = NLDTM.
              label = 'Datetime';
```

Note: **DATETIME** is the starting date and time of interval data. If the raw data provides the ending date and time, you should subtract the interval duration from the date and time in order to calculate the starting date and time.

- **LSTPDATE**

In order to take advantage of the duplicate checking process that is available in SAS IT Resource Management aggregations, include a column called **LSTPDATE** in your staging tables. This column should be set to the current date and time: **LSTPDATE = datetime();**

Its attributes are as follows:

```
attrib LSTPDATE length = 8
              format = NLDTM.
              label = 'LastProcessDate';
```

- **DURATION**

Interval-type data is data that is captured or snapshot at given intervals. For interval type data, supplied adapters include a field called **DURATION**. This

field is set to the duration of each interval in seconds. It is useful as the weight variable in an aggregation when creating weighted statistics. Its attributes are as follows:

```
attrib DURATION length = 8
        format = TIME8.
        label = 'Duration';
```

Assign an appropriate format and label to your other columns. (You can also do this at a later time in the SAS IT Resource Management client in steps 9.k through 9.m.)

Note: It is not necessary to create aggregation variables such as Hour, Shift, WeekDate, and so on. You can add these columns in step 9.l.

6. Run the program to create the staged tables.
7. On the SAS IT Resource Management client, register your staged tables.

How to register your staged tables.

- a. On the SAS IT Resource Management client, right-click the staging library that you created in steps 3c through 3g, and select **Register Tables**. The Register Tables wizard opens.
 - b. On the first page of the wizard, you do not need to make any changes. Your library is listed on that page.
 - c. The tables that you created in step 6 are listed. Click **Select All Tables** to register all the tables in metadata.
 - d. Continue to follow the prompts of the wizard to complete the registration.
8. If the input to the User-Written Staging transformation is a SAS data set or another type of database table, then create the metadata for it.

TIP One method that you can use to create the necessary metadata for your input is to create a library in metadata that points to your source data. Then select **File** ⇒ **Register Tables** to register the tables in metadata.

9. Create a staging job to populate the staging tables.

How to create a staging job.

- a. Right-click the **Staging** folder in your IT data mart and select **New** ⇒ **Job**.
- b. Give the job a name such as Staging.
- c. From the Transformations tree, expand **ITRM Transformations**. Then drag the User-Written Staging transformation onto the job that you just created in the **Diagram** tab of the Job Editor window.
- d. Drag the staging tables from the **Staging** folder in your IT data mart onto the job.
- e. Connect the User-Written Staging transformation to each of the staging tables so that the tables are outputs of the transformation.
- f. If you created input metadata in step 8, then drag the input table to the job. Connect the table as input to the User-Written Staging transformation.
- g. Right-click the User-Written Staging transformation and select **Properties**. From the drop-down list, select the **Staging Parameters** tab. The Staging page appears.

- a. On the Staging page, specify the path to the raw data that you want to process. A macro variable called Rawdata will be created in the generated code with this value.

On this page, you can also specify duplicate checking and future data options.

- b. Select the **Duplicate Checking** option in the left panel. The Duplicate Checking page appears. It provides a set of parameters that will be used if you decide to implement the duplicate checking process. For information about how to implement duplicate checking of your data, see [“Duplicate Checking” on page 242](#).
- c. Select the **User-Written** option in the left panel. The User-Written page appears. The first parameter on this page enables you to specify the path to your SAS staging program. A value for this parameter is required.

The other parameters enable you to specify what macros and macro variables you want the transformation to generate for your SAS program.

Note: For information about these properties, see [Step 8 of the “Create a User-Written Staging Transformation” topic on page 253](#).

TIP You can click the **Reset to defaults** option at the top of the page to reset all the parameters in that subgroup to the values that were established when SAS IT Resource Management was installed. Click the **Reset** option next to any one of these parameters to reset it to the value that was established when SAS IT Resource Management was installed. Also note that **Reset to defaults** does not reset a parameter that is dependent on another parameter if the parent parameter does not change.

- d. Standard SAS IT Resource Management formulas are provided to compute columns such as WEEKDATE and SHIFT. If you want to use any of these formulas, double-click the tables in the job that is opened in the **Diagram** tab of the Job Editor window. Then select the **Columns** tab.
- e. On the **Columns** tab, you can add the standard SAS IT Resource Management date columns and their formulas.

The following table lists the standard columns.

Table 11.1 Table of Standard SAS IT Resource Management Columns

Column Name		
Length		
Format	ITRM Formula	Description
DATETIMEFIFTEENMIN Length: 8 Format: NLDATM.	DatetimeFifteenMinute	DatetimeFifteenMinute
DATETIMEFIVMIN Length: 8 Format: NLDATM.	DatetimeFiveMinute	DatetimeFiveMinute
DATETIMEONEMIN Length: 8 Format: NLDATM.	DatetimeOneMinute	DatetimeOneMinute

Column Name		
Length		
Format	ITRM Formula	Description
DATETIMETENMIN Length: 8 Format: NLDATM.	DatetimeTenMinute	DatetimeTenMinute
DATETIMETHIRTYMIN Length: 8 Format: NLDATM.	DatetimeThirtyMinute	DatetimeThirtyMinute
DATETIMETWENTYMIN Length: 8 Format: NLDATM.	DatetimeTwentyMinute	DatetimeTwentyMinute
DAYDATE Length: 8 Format: NLDATE10.	DayDate	DayDate
DAYOFMONTH Length: 8 Format: 2.	DayOfMonth	DayOfMonth
DAYOFWEEK Length: 8 Format: WKD.	DayOfWeek	DayOfWeek
HOUR Length: 8 Format: 2.	Hour	Hour
MONTHDATE Length: 8 Format: NLDATE10.	MonthDate	MonthDate
SHIFT Length: \$1 Format: \$CHAR1.	Shift	Shift
TIME Length: 8 Format: TOD5.	Time	Time
WEEKDATE Length: 8 Format: NLDATE10.	WeekDate	WeekDate

Column Name		
Length		
Format	ITRM Formula	Description
YEARDATE	YearDate	YearDate
Length: 8		
Format: NLDATE10.		

TIP It is easier to add these computed columns to the metadata using the SAS IT Resource Management client than to include them in the physical table that you used to register the table. If you include them when you register the table, then they will all be created as “Data Columns” instead of “Computed Columns.” In that case, you must delete the columns and then recreate them as “Computed Columns.”

From the **Columns** tab, you can import columns from another source or add new columns.

Import columns to your staged table.

- a. Click the Import Columns icon.

TIP The Import Columns icon is next to the **New** button.

- b. In the Import Columns dialog box, expand **Adapters**. Then drill down to any standard adapter template table.

TIP One of the SAR adapter template tables is a good choice because these tables are small.

- c. Choose the columns that you want to include in your staging tables. For information about available columns, see the preceding list of columns.
- d. Click **OK**. The columns are added to your table with the appropriate attributes and formulas.

Add new columns to your staged table.

- e. Click the **New** button.
- f. You can add data and computed columns to your staged table. For information, see [“Columns in Staged Tables” on page 124](#).
- g. If you add computed columns, the automatically generated macro called `_ITMS_computedColumns` can be used in your SAS program to calculate the values of these columns.

CAUTION:

When creating a computed column, do not base it on another computed column, because the code might not calculate the columns in the correct order. For example, if you are computing column A as `1+1`, and you want another computed column called B that is `A+5`, then the expression for B should be `(1+1)+5`, not `A+5`.

Expressions can be simple or complex.

- A simple expression is only the content on the right side of an equation without a semicolon.
- A complex expression contains multiple statements including the necessary semicolons. The following code is an example of a complex expression.

```
if 1/A (A ne 0) then B=1/A; else B=0;
```

This code prevents division by zero, which causes errors when the code runs.

10. Save the job and deploy the SAS code. To do so, right-click the job and select **Scheduling** ⇒ **Deploy**. For more information, see [“Deploy a Job for Scheduling” on page 274](#).
11. To make your SAS program more robust, you might want to use some of the macro variables and macros that are provided by SAS IT Resource Management. These macros and macro variables enable your program to handle changes in metadata without changing the program. The following topics discuss some of the changes that you can make to your program.
 - [“Duplicate Checking” on page 242](#)
 - [“Future Checking” on page 243](#)
 - [“Error Checking” on page 244](#)
 - [“LIBNAME Statements” on page 244](#)
 - [“Formats” on page 244](#)
 - [“Helper Macros and Macro Variables” on page 244](#)

For examples of how macros and macro variables can be used, see [“Example of Modifying the Code ” on page 245](#).

Duplicate Checking

A set of macros is provided to check for duplicate data. These macros are used in all the supplied adapters to handle reading and saving duplicate data. If your SAS program finds data in the raw data that has already been read, then the data is handled based on the setting of the **Duplicate checking** option (on the Duplicate Checking section of the **Staging Parameters** tab):

- Discard—delete the duplicate data from the input stream
- Terminate—stop processing
- Force—ignore any duplicate data and continue processing
- Inactive—do not perform any duplicate data checking. This option permits the duplicate data to be processed.

You can use some of the helper macros from the User-Written Staging transformation as well as the standard duplicate-data checking macros to handle duplicate-data checking. For more information about duplicate-data checking macros, see [“Duplicate-Data-Checking Macros” on page 394](#). To implement duplicate checking in your code, perform the following steps.

1. In the DATA statement of the DATA step that reads the raw data, add this macro call:

```
%RMDUPDSN (SOURCE=USR)
```

This creates a work data set with duplicate information. For example, suppose that the following is the current data statement:

```
data stage.stageTable1
    stage.stageTable2;
```

You need to change it to this:

```
data stage.stageTable1
      stage.stageTable2
      %RMDUPDSN(SOURCE=USR) ;
```

2. At the beginning of the DATA step, add this macro call:

```
%_ITMS_dupInit;
```

This initializes the duplicate-data checking process.

3. On the SET or INFILE statement, be sure to include the END= option. This option should point to the ENDFILE variable specified on the Duplicate Checking section of the **Staging Parameters** tab of the User-Written Staging Properties dialog box.

4. In the DATA step, you also need to add this macro call:

```
%_ITMS_dupCheck;
```

This macro runs the %RMDUPCHK macro, which performs the duplicate-data checking. The parameters for the %RMDUPCHK macro are specified on the Duplicate Checking section of the **Staging Parameters** tab of the User-Written Staging Properties dialog box. You can call this macro multiple times in your code if the TIMESTMP column is set in different places.

TIP Put this macro call in the DATA step at any point after the TIMESTMP column has a value. TIMESTMP is a parameter in the %RMDUPCHK macro that specifies the date and time column name. (Typically, this column is DATETIME, but you can call it anything you want.) Specify the column to use on the **Staging Parameters** tab of the User-Written Staging Properties dialog box.

5. Following the DATA step that ran the _ITMS_dupCheck macro, add this macro call:

```
%RMDUPUPD;
```

This macro completes the duplicate checking process by adding the new date and time stamps from the data that was just read to the duplicate checking tables that are kept in the staging library.

Future Checking

A set of macros is provided to check for data that has a date and time stamp that is in the future. (This situation probably means that the data is invalid). These macros assume that your TIMESTMP variable (as described in the previous duplicate checking section) is called DATETIME.

To implement future checking, you can use a User-Written Staging transformation helper macro as well as the standard future processing macros:

- In the DATA step that you use to read the raw data, add this macro call after the DATETIME column has been set:

```
%RMSETFTR;
```

Note: You can call this macro multiple times if the DATETIME column is set in multiple places.

- After the DATA step that ran the %RMSETFTR macro, add this macro call:

```
%_ITMS_futureData;
```

This macro invokes the %RMFUTURE macro for all your staging tables.

Error Checking

If you do error checking in your program, you can take advantage of the abort routine that is provided at the end of the generated code. If you set the macro variable `&trans_rc` to a value greater than 4, then the program will be cancelled and set its return code to 8 (`abort cancel 8;`).

To set the `&trans_rc` and the `&job_rc` error flags outside of a PROC step or DATA step, you can use the `%RCSET(ERROR)` macro. To set those flags inside a DATA step, you can use the `%RCSETDS(ERROR)` macro.

The error parameter for these macros is the error code that you want to specify. For example, you can use the following code to set a return code of 8.

```
%rcSet (8) ;
```

You can use the following code to set the return code to the error code from the proceeding DATA step or PROC step.

```
%rcSet (&syserr) . ;
```

The macro definitions for `%RCSET` and `%rRCSETDS` can be found at the beginning of the generated code.

LIBNAME Statements

All the LIBNAME statements for the source and target tables will be generated for you. Therefore you do not need to include LIBNAME code in the user-written code. However, if you are using tables that are not in the metadata as source or target tables, then you should issue your own LIBNAME statements.

Formats

To access standard SAS IT Resource Management formats, set the FMRSEARCH option as follows:

```
options fmtsearch=(admin.formats sashelp.itms_formats);
```

Your own formats can be put in the ADMIN.FORMATS catalog or in your own catalog. If you use a different catalog, you need a LIBNAME statement for that library. You then need to set the FMTSEARCH option to point to your formats. For example:

```
options fmtsearch=(admin.formats mylib.formats sashelp.itms_formats);
```

The NOFMterr option will be set for you so that processing does not fail if there is a missing format.

Helper Macros and Macro Variables

A set of macros and macro variables can be created for you in the generated code. On the **Staging Parameters** tab of the User-Written Staging Properties dialog box, select **User-Written** to access these options. For information about the macro variables and macros that are available, see [“Macro Variables” on page 260](#) and [“Macros” on page 260](#).

TIP Set all the options on the macro variables to **Yes**, and then generate the code. You can then see what macros and macro variables are available and determine whether you want to use them. You can then turn off the macros and macro variables that you do not need.

On the User-Written section of the **Staging Parameters** tab of the User-Written Staging Properties dialog box, you can also specify the creation of the Tabinfo and Varinfo tables. For information about these tables, see [“Tabinfo and Varinfo” on page 261](#).

Example of Modifying the Code

The following two examples show a simple user-written staging program before it is modified, and then after it is modified to include the helper macros.

Example Code 11.1 Original Program

```

/* ***** */
/* Name: UserWrittenStaging */
/* ----- */
/* Description: */
/* Reads the rawdata from the rawdata.txt file that is in */
/* this format: */
/* ddMONyyyy:hh:mm machineName metric1 metric2 */
/* ddMONyyyy:hh:mm machineName metric1 metric2 */
/* <etc.> */
/* ***** */

libname staging 'c:\UserWritten\Staging';
data staging.metric1 (keep = datetime machine
                    metric1 lstpdate)
    staging.metric2 (keep = datetime machine
                    metric3 lstpdate);
    attrib datetime label = "DateTime" format = datetime18.;
    attrib machine label = "Machine" length = $20;
    attrib metric1 label = "Metric1" format = NLNUM16.;
    attrib metric3 label = "Metric2*100" format = NLNUM16.;
    attrib lstpdate label = "LastProcessDate" format = NLDATM.;
    attrib metric2 format = NLNUM16.;
    lstpdate = datetime(); /* needed for aggregation */
    infile "c:\UserWritten\rawdata.txt" missover;
    input datetime datetime15. machine metric1 metric2;
    metric3 = metric2 * 100;
    output staging.metric1
        staging.metric2;

run;

```

Here is what the program might look after adding the helper macros.

Example Code 11.2 Modified Program

```

/* ***** */
/* Name: UserWrittenStaging */
/* ----- */
/* Description: */
/* Reads the rawdata from the rawdata.txt file that is in */
/* this format: */
/* ddMONyyyy:hh:mm machineName metric1 metric2 */
/* ddMONyyyy:hh:mm machineName metric1 metric2 */
/* <etc.> */
/* ***** */
/* ***PUT CODE INTO A MACRO IN ORDER TO USE %DO STATEMENTS*** */
%macro readRawdata;
/* libname staging 'c:\UserWritten\Staging'; */
data
/* ***GENERATE TABLE NAMES WITH OPTIONS AND KEEP LISTS*** */

```

```

        %do i = 1 %to &numTargets;
            &&target&i (%_ITMS_tableOptions(targetTableNum=&i,
                                                generateTableName=NO)
                        keep = %_ITMS_columnList(targetTableNum=&i))
        %end;
/* ***DUPLICATE CHECKING OUTPUT TABLE*** */
    %RMDUPDSN(SOURCE=USR)
/*      staging.metric1 (keep = datetime machine */
/*                          metric1 lstpdate) */
/*      staging.metric2 (keep = datetime machine */
/*                          metric2 metric3 lstpdate) */
    ;
/* ***INITIALIZE DUPLICATE CHECKING*** */
    %_ITMS_dupInit;
/* ***GENERATE ATTRIB STATEMENTS FOR COLUMNS IN OUTPUT TABLES*** */
    %_ITMS_attrib();
/*      attrib datetime label = "DateTime" format = datetime18.; */
/*      attrib machine label = "Machine" length = $20; */
/*      attrib metric1 label = "Metric1" format = NLNUM16.; */
/*      attrib metric3 label = "Metric2*100" format = NLNUM16.; */
/*      attrib lstpdate label = "LastProcessDate" format = NLDATM.; */
/*      attrib metric2 format = NLNUM16.; */
/*      lstpdate = datetime(); /* needed for aggregation */
/*      infile "c:\UserWritten\rawdata.txt" missover; */
/*      ***USE &RAWDATA AND INCLUDE AN END-OF-FILE FLAG FOR*** */
/*      ***DUPLICATE CHECKING*** */
/*      infile "&rawdata" end = _eof missover; */
/*      input datetime datetime15. machine metric1 metric2; */
/*      ***RUN DUPLICATE CHECKING*** */
/*      %_ITMS_dupCheck; */
/*      ***RUN FUTURE DATA CHECKING*** */
/*      %RMSETFTR; */
/*      ***CREATE COMPUTED COLUMNS*** */
/*      %_ITMS_computedColumns; */
/*      metric3 = metric2 * 100; */
/*      ***OUTPUT DATA TO TARGET TABLES*** */
    output
        %do i = 1 %to &numTargets
            &&target&i
        %end;
    ;
/*      output staging.metric1 */
/*      staging.metric2; */
run;

/* ***UPDATE DUPLICATE CHECKING FILES*** */
    %RMDUPUPD;
/* ***HANDLE FUTURE DATA*** */
    %_ITMS_futureData;
%mend readRawdata;
%readRawdata;

```

Aggregating Data Created by User-Written Staging Transformations

After the data is staged, it is ready to be aggregated. The data in the staged tables that you created can be rolled up into aggregation tables. These tables are then available for analysis.

Creating and working with aggregations is described in detail in chapter 8, “Aggregating the Data.” For information, see [“About Aggregation Tables” on page 132](#).

Use the following tips to keep your user-written adapter consistent with the supplied adapters.

- In the IT data mart that you created for your user-written adapter, create a folder called **Domain Categories**. Within that folder, create one or more folders for your domain categories. (A domain category is a way of grouping different types of data such as System, Disk, or Network data.) Specify a name for these categories that makes sense for your application.

Many of the supplied adapters have a domain category for each of the staged tables that were created. If you do not have domain categories, create a single folder that has a descriptive name.

- Create an aggregation job for each domain category. Specify a job name using this format: *<domain category> Aggregation*.
- In the supplied adapters, the data in each staged table is used as input to multiple aggregations. The data is typically summarized into different time periods. Most supplied adapters include the following time period aggregations.

Table 11.2 Supplied Time Period Aggregations

Aggregation Table Prefix	Date and Time Class Columns
Day	DAYDATE
DayHour	DAYDATE HOUR
DayShift	DAYDATE SHIFT
DayShiftHour	DAYDATE SHIFT HOUR
Week	WEEKDATE
WeekHour	WEEKDATE HOUR
WeekShift	WEEKDATE SHIFT
WeekShiftHour	WEEKDATE SHIFT HOUR
Month	MONTHDATE
MonthHour	MONTHDATE HOUR
MonthShift	MONTHDATE SHIFT
MonthShiftHour	MONTHDATE SHIFT HOUR

In addition, many adapters supply the following aggregations:

Table 11.3 Other Supplied Aggregations

Aggregation Table Prefix	Description
Detail	This data is not summarized. It is left at the same level as the staging data.
KeyMetrics	Generally, this data is summarized into the DAYDATE and HOUR level. It also includes a limited number of metrics. This table usually keeps data for a longer period than the other aggregations, so that it can be used in trending and forecasting.

TIP For the best results, use these standard prefixes on your aggregation tables. The tables will then conform to the standard that is used for the supplied adapters.

- Consider using rank columns if needed.

Rank columns can be used to create “Top N” reports and ranked listings. They also provide a way to filter the data. With rank columns, you can filter the data so that only data from a certain time period (relative to the data that has been collected) is processed in the reports. For example, if you want the data for yesterday, you can select all values that have a value of 1 for the rank variable called DayDateRank01. Some standard data and time rank columns are listed in the following table.

Table 11.4 Standard Data and Time Rank Columns

Summarization Level of Aggregation Tables	Staging Column	Rank Column	Rank Column Label
For all aggregation tables that are summarized at the Day level and more granular	DayDate	DayDateRank01	DayDateDescRank
For all aggregation tables that are summarized at the Week level and more granular	WeekDate	WeekDateRank01	WeekDateDescRank
For all aggregation tables that are summarized at the Month level and more granular	MonthDate	MonthDateRank01	MonthDateDescRank

- Consider using join columns if needed.

Join columns are usually rank columns that are copied from another aggregation table. These columns can be used to filter the data for Top 10 reports. For example, a rank column in the DayHour aggregation table is the ranking of that column for each day and hour.

You might want to filter the data by the top 10 for the day, even if you are showing the data by hour. If so, you need to create a join column in the DayHour table in

order to get the ranking column from the Day table. By default, the supplied aggregations have join columns for all the non-date and time ranking columns. However, in the user-written adapter, you can create only the ones that you will need.

The naming convention for join columns is the following: *X_rankColumn*, where *rankColumn* is the name of the rank column from the original table and *X* is one of the following values:

- *D* if the rank column is from the Day table
- *W* if the rank column is from the Week table
- *M* if the rank column is from the Month table

Creating Information Maps for User-Written Adapters

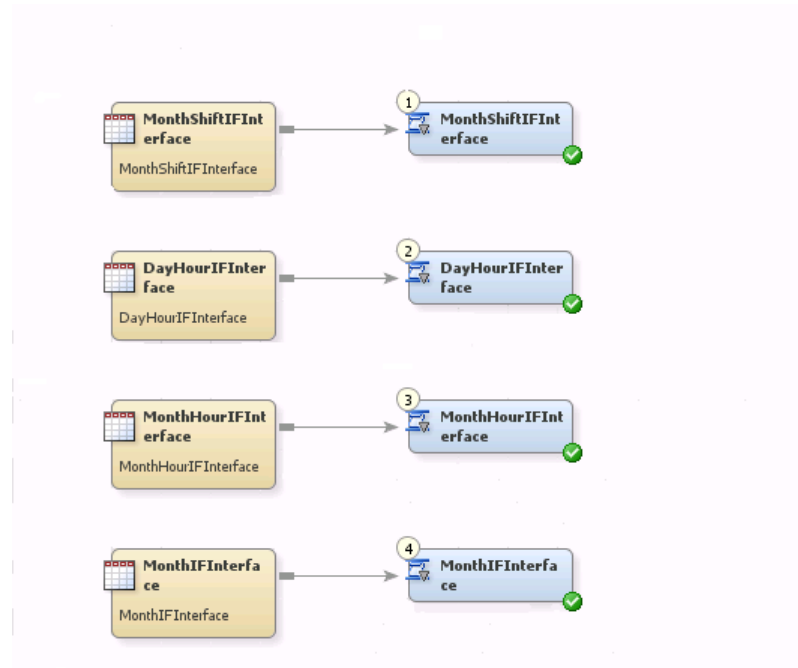
The supplied adapters contain jobs that, when run, create information maps for each of the aggregation tables that were generated. You can also create information map jobs for your user-written aggregations. An information map enables you to access your data with SAS Web Report Studio, SAS Enterprise Guide, and other SAS reporting tools.

Note: In SAS Enterprise Guide, you can also access the data directly from the SAS data tables.

Creating and working with information maps is described in detail in chapter 10, “Information Maps.” For more information, see [“About Information Maps” on page 208](#).

The following tips can be used to keep your user-written adapter consistent with the supplied adapters.

- Create an information map job for each domain category. Specify a name for the job such as *<domain category>Information Map*.
- Each job should have an Information Map transformation for each aggregation in the domain category. The following diagram shows the resulting Information Map transformations for the job.

Display 11.1 Information Map Transformations

- Use the name of the aggregation table as the name of your information map.
- Information map filters are helpful when using information maps for reporting. Some standard filters that are used in the supplied adapters are shown in the following tables. The filters are listed by the columns on which the aggregation table is summarized: DayDate, WeekDate, and MonthDate.

Table 11.5 Standard Filters Used for Supplied Adapters Based on the DayDate Column

Filter Name	Filter Expression and Description
LastDay	Expression: <<aggTableNameDAYDATERANK01>> = 1 Description: DayDateDescRank 1
Last3Days	Expression: <<aggTableNameDAYDATERANK01>> between 1 and 3 Description: DayDateDescRank 1 – 3
RollingWeek	Expression: <<aggTableNameDAYDATERANK01>> between 1 and 7 Description: DayDateDescRank 1 - 7
RollingMonth	Expression: <<aggTableNameDAYDATERANK01>> between 1 and 31 Description: DayDateDescRank 1 - 31
Last14Days	Expression: <<aggTableNameDAYDATERANK01>> between 1 and 14 Description: DayDateDescRank 1 - 14
Last90Days	Expression: <<aggTableNameDAYDATERANK01>> between 1 and 90 Description: DayDateDescRank 1 - 90

Table 11.6 Standard Filters Used for Supplied Adapters Based on the WeekDate Column

Filter Name	Filter Expression and Description
ThisWeek	Expression: <<aggTablenameWEEKDATERANK01>> = 1 Description: WeekDateDescRank 1
LastWeek	Expression: <<aggTablenameWEEKDATERANK01>> = 2 Description: WeekDateDescRank 2
Last2Weeks	Expression: <<aggTablenameWEEKDATERANK01>> between 2 and 3 Description: WeekDateDescRank 2 - 3
Last4Weeks	Expression: <<aggTablenameWEEKDATERANK01>> between 2 and 5 Description: WeekDateDescRank 2 - 5
Last12Weeks	Expression: <<aggTablenameWEEKDATERANK01>> between 2 and 13 Description: WeekDateDescRank 2 - 13

Table 11.7 Standard Filters Used for Supplied Adapters Based on the MonthDate Column

Filter Name	Filter Expression and Description
ThisMonth	Expression: <<aggTableName.MONTHDATERANK01>> = 1 Description: MonthDateDescRank 1
LastMonth	Expression: <<aggTableName.MONTHDATERANK01>> = 2 Description: MonthDateDescRank 2
Last2Months	Expression: <<aggTableName.MONTHDATERANK01>> = 2 - 3 Description: MonthDateDescRank 2 - 3
Last4Months	Expression: <<aggTableName.MONTHDATERANK01>> = 2 - 4 Description: MonthDateDescRank 2 - 4
Last13Months	Expression: <<aggTableName.MONTHDATERANK01>> = 2 - 14 Description: MonthDateDescRank 2 - 14

You can also make filters for the other ranked and joined columns in your data. For the supplied adapters, there are filters for all the rank and join columns in the aggregation. These filters enable you to create top-ten reports. To generate this type of report, specify “between 1 and 10” in the filter expression.

Reporting on Data Created with User-Written Staging Transformations

To create reports that can be displayed in the Gallery Manager, you must use SAS Enterprise Guide. For information about creating and using reports within the SAS IT Resource Management 3.2 framework, see the *SAS IT Resource Management 3.2: Reporting Guide*. This document can be accessed from the <http://support.sas.com/itrm> Web site.

Working with User-Written Staging Transformations

Create a User-Written Staging Transformation

To create a User-Written Staging transformation, perform the following steps:

1. From the IT Data Marts tree of SAS Data Integration Studio, double-click the job that is to contain the User-Written Staging transformation. The job opens on the **Diagram** tab of the Job Editor window.
2. From the Transformations tree of SAS Data Integration Studio, expand the **ITRM Transformations** folder.
3. Scroll down to the **User Written Staging** entry. Drag and drop it onto the job on the **Diagram** tab of the Job Editor window.
4. (Optional) Attach a source table on the left side of the transformation.

To create a table, select **File** ⇒ **New** ⇒ **Table** to open the New Table wizard.

Respond to the prompts as needed and click **Finish**. To attach the new table to the transformation, navigate to the location of the table in the tree where you stored its metadata. Drag and drop the table onto the process flow diagram and attach the table to the transformation.

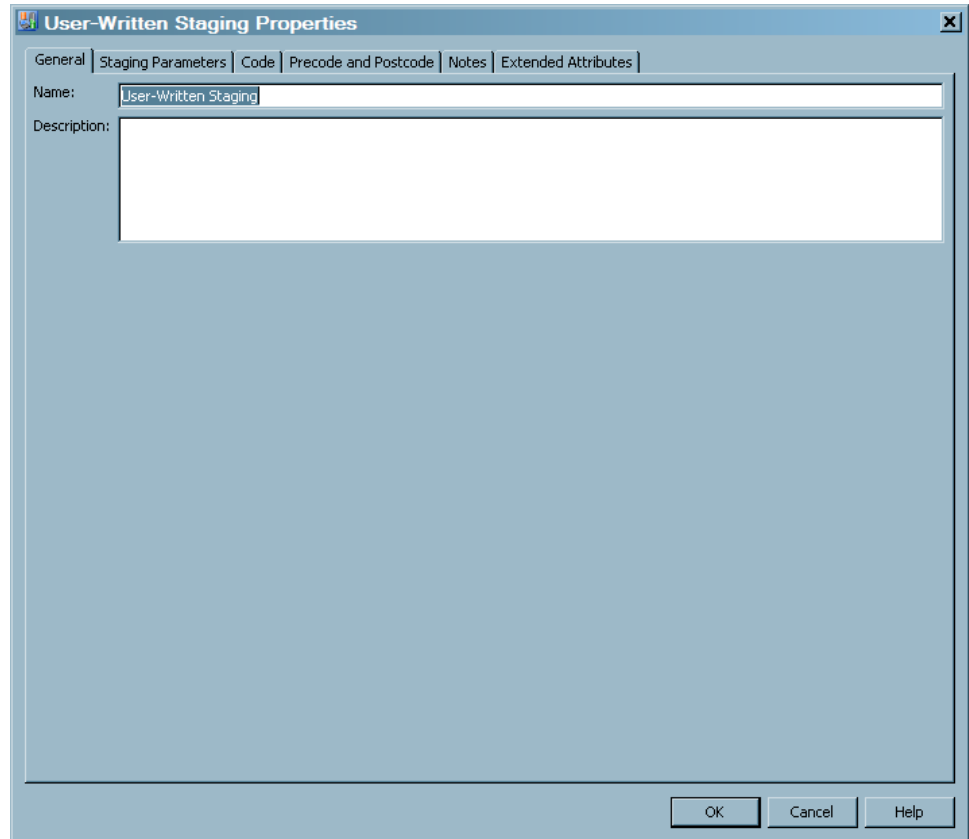
5. Attach a target table to the transformation. (Multiple target tables can be attached to the transformation.)

TIP If you cannot attach the table to the transformation, right-click the transformation and from the menu that is displayed, select **Connections**. In the User Written Staging Connections dialog box that opens, the **Output Node** field of the target table should have a value. If the field does not have a value (or displays **<none>**), double-click the ellipsis (...) in the **Output Node** field and select the table to which you want to connect. Then click **OK**. Click **OK** again to return to the process flow diagram and connect the target table to the transformation.

6. Right-click the User-Written Staging transformation. From the drop-down menu, select **Properties**.

The **General** tab of the User-Written Staging Properties dialog box appears. The **Name** field displays the default value of the transformation, **User-Written Staging**.

Display 11.2 General Tab of the User-Written Staging Transformation

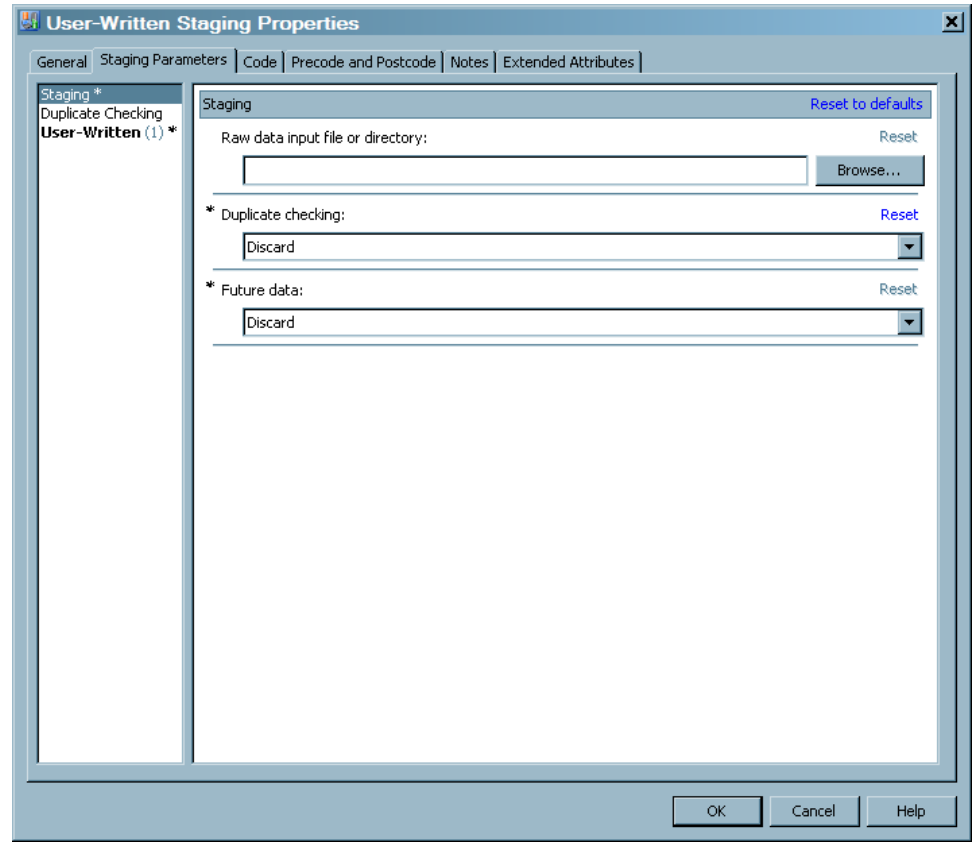


7. In the **Name** field on the **General** tab, enter the name of the transformation if you do not want to retain the default name. The name of the transformation can be no more than 60 characters. It must be unique within the job that contains it. This field cannot be blank.

The new name of the transformation appears in the box that represents the User-Written Staging transformation in process flow diagram.

In the **Description** field, you can enter up to 200 characters of information about the transformation. A value is optional in this field.

8. On the **Staging Parameters** tab, you can customize the staging transformation.
 - a. In the left panel of the tab, select **Staging**.

Display 11.3 Standard Staging Parameters

You can specify the following information:

- In the **Raw data input file or directory** field, enter the location of the file or directory that contains the raw data that is input to your staging transformation.

Click **Browse** to open the Select the File or Directory dialog box. In the **Look in** field, the drop-down list displays the files and directories that are available in the metadata repository to which you are connected. From the drop-down list that is displayed, select the file or directory that contains the raw data that is input to your user-written staging transformation.

For information about the Raw data input file or directory parameter, see [“Choose Raw Data Input Type” on page 357](#).

- In the **Duplicate Checking** field, the drop-down list displays the options that are available for handling duplicate-data checking. These options are **Discard**, **Force**, **Inactive**, and **Terminate**. Select the option that you want to use, or use the default value, **Discard**.

For information about the Duplicate Checking parameter, see [“Duplicate Checking” on page 357](#).

- In the **Future Data** field, the drop-down list displays the options that are available for handling future data. These options are **Accept**, **Discard**, and **Terminate**. Select the option that you want to use, or use the default value, **Discard**.

For information about the Future Data parameter, see [“Future Data” on page 359](#).

Note: Click the **Reset to defaults** option at the top of the tab to reset all the Staging parameters to the values that were established when SAS IT Resource Management was installed. Click the **Reset** option next to any one of these parameters to reset it to the value that was established when SAS IT Resource Management was installed. Also note that **Reset to defaults** does not reset a parameter that is dependent on another parameter if the parent parameter does not change.

- b. In the left panel of the tab, select **Duplicate Checking**. The following dialog box appears.

Display 11.4 Duplicate Checking Staging Parameters

The screenshot shows the 'User-Written Staging Properties' dialog box with the 'Duplicate Checking' tab selected. The left pane lists 'Staging*', 'Duplicate Checking', and 'User-Written (1)*'. The main area contains several input fields, each with a 'Reset' button to its right. The fields are: 'IDVAR variable for %RMDUPCHK macro:' (value: MACHINE), 'TIMESTMP variable for %RMDUPCHK macro:' (value: DATETIME), 'ENDFILE variable for %RMDUPCHK macro:' (value: _eof), 'INT value for %RMDUPCHK macro:' (value: 00:29), 'RANGES value for %RMDUPCHK macro:' (value: 20), 'SYSTEMS value for %RMDUPCHK macro:' (value: 3), and 'KEEP value for %RMDUPCHK macro:' (value: 9). At the top right of the main area is a 'Reset to defaults' link. At the bottom right are 'OK', 'Cancel', and 'Help' buttons.

You can specify the following information, which pertains to the parameters and options that affect how duplicate data checking is performed:

- In the **IDVAR for the %RMDUPCHK macro** field, you can specify the name of the SAS variable that identifies the system or machine that generated the input data.

The default value is **MACHINE**.

- In the **TIMESTMP variable for the %RMDUPCHK macro** field, you can specify the name of the SAS variable that contains the datetime stamp that uniquely identifies the time of the event or interval that is being recorded.

The default value is **DATETIME**.

- In the **ENDFILE variable for the %RMDUPCHK macro** field, you can specify the name of the SAS variable that is used as the END= keyword for the SAS INFILE statement that reads the raw data.

The default value is **_eof**.

- In the **INT value for the %RMDUPCHK macro** field, you can specify the maximum time gap (or interval) that is to be allowed between the timestamps on any two consecutive records from the same system or machine. If the interval between the timestamp values exceeds the value of this parameter, then an observation with the new time range is created in the control data set. This is referred to as a gap in the data.

The value for this parameter must be provided in the format hh:mm, where hh represents hours and mm represents minutes. For example, to specify an interval of 14 minutes, use INT=0:14. To specify an interval of 1 hour and 29 minutes, use INT=1:29.

The default value is **00:29**.

TIP For best results, specify the interval to be a little larger than the natural interval of the raw data. For example, if the data is in 10-minute intervals, set INT=00:12. If the data is hourly, set INT=01:02.

- In the **KEEP value for the %RMDUPCHK macro** field, you can specify the number of weeks for which control data are to be kept. Because this value represents the number of Sundays between two dates, a value of 2 (the default) results in a maximum retention period of 20 days. This value must be an integer.

The default value is **9**.

Note: Click the **Reset to defaults** option at the top of the tab to reset all the Duplicate Checking parameters to the values that were established when SAS IT Resource Management was installed. Click the **Reset** option next to any one of these parameters to reset it to the value that was established when SAS IT Resource Management was installed. Also note that **Reset to defaults** does not reset a parameter that is dependent on another parameter if the parent parameter does not change.

- c. In the left panel of the tab, select **User-Written**. The following dialog box appears.

Display 11.5 User-Written Staging Parameters

The screenshot shows the 'User-Written Staging Properties' dialog box with the 'Staging Parameters' tab selected. The left pane shows 'Staging *', 'Duplicate Checking', and 'User-Written (1) *'. The right pane contains several settings for the 'User-Written' staging transformation, each with a 'Reset' button.

Setting	Value	Reset
* User-written code:	Select the file that contains the user-written SAS code. [Text Field] [Browse...]	Reset
* Generate a macro variable with input table name:	%input [Yes] [v]	Reset
* Generate macro variables with target table names:	%target1-%targetn [Yes] [v]	Reset
* Generate a macro with target table options:	%_ITMS_tableOptions(targetTableName=, generateTableName=, generateIndexes=) [Yes] [v]	Reset
* Generate Tabinfo and Varinfo tables:	Temporary SAS tables containing information about target tables and variables [No] [v]	Reset
* Generate a macro to create target table ATTRIB statements:	%_ITMS_attrib(targetTableName=) [Yes] [v]	Reset
* Generate a macro to create computed columns:	%_ITMS_computedColumns(targetTableName=) [Yes] [v]	Reset
* Generate a macro to create column lists:	%_ITMS_columnList(targetTableName=, delimiter=, sqlAlias=, quotes=NO, columnType=ALL) [Yes] [v]	Reset
* Generate a macro to assign columns to missing:	%_ITMS_assignToMissing(targetTableName=) [Yes] [v]	Reset

Buttons: OK, Cancel, Help

You can specify the following information:

- In the **User written code** field, enter the location where your staging code is stored.

Click **Browse** to open the Select a File dialog box.

- In the **Generate a macro variable with input table name** field, you can choose to generate a macro variable that contains the full table name (**libref.table**) of the input table. From the drop-down list, select **Yes** if you want to generate the macro variable. Otherwise, select **No**.
- In the **Generate macro variables with target table names** field, you can choose to generate macro variables that contain the full table name (**libref.table**) of the target tables. From the drop-down list, select **Yes** if you want to generate the macro variables. Otherwise, select **No**.
- In the **Generate a macro with target table options** field, you can choose to generate a macro that contains the data set options for each target table. From the drop-down list, select **Yes** if you want to generate the macro. Otherwise, select **No**.

- In the **Generate Tabinfo and Varinfo tables** field, you can choose to generate code that creates the following tables:
 - **tabinfo**, which contains all the target tables
 - **varinfo**, which contains all the columns in all the target tables

From the drop-down list, select **Yes** if you want to generate the code. Otherwise, select **No**.
- In the **Generate a macro to create target table ATTRIB statements** field, you can choose to generate a macro that contains **attrib** statements for each column in the target tables. From the drop-down list, select **Yes** if you want to generate the macro. Otherwise, select **No**.
- In the **Generate a macro to create computed columns** field, you can choose to generate a macro that contains the computations for any computed columns in the target tables. From the drop-down list, select **Yes** if you want to generate the macro. Otherwise, select **No**.
- In the **Generate a macro to create column lists** field, you can choose to generate a macro that creates a list of all the columns in the table that can be used as a keep list, a retain list, or any other list of columns that is needed. From the drop-down list, select **Yes** if you want to generate the macro. Otherwise, select **No**.
- In the **Generate a macro to assign columns to missing** field, you can choose to generate a macro that contains assignment statements to set all columns in the target table to missing. From the drop-down list, select **Yes** if you want to generate the macro. Otherwise, select **No**.

Note: Click the **Reset to defaults** option at the top of the tab to reset all the User Written parameters to the values that were established when SAS IT Resource Management was installed. Click the **Reset** option next to any one of these parameters to reset it to the value that was established when SAS IT Resource Management was installed. Also note that **Reset to defaults** does not reset a parameter that is dependent on another parameter if the parent parameter does not change.

9. The remaining tabs on the User-Written Staging Properties dialog box consist of standard SAS Data Integration Studio properties. For information about how to use these properties, click **Help** on the tab.
10. Click **OK** to return to the process flow diagram.
11. Click **File** ⇒ **Save** to save the job.

Delete a User-Written Staging Transformation from a Job

To delete a User-Written Staging transformation from a job, perform the following steps:

1. From the IT Data Marts tree of SAS Data Integration Studio, double-click the job that contains the User-Written Staging transformation. The job opens on the **Diagram** tab of the Job Editor window.
2. Right-click the box that represents the User-Written Staging transformation.
3. From the drop-down menu, select **Delete**.
4. Click **File** ⇒ **Save** to save the job.

Modify a User-Written Staging Transformation

To modify a User-Written Staging transformation, perform the following steps:

1. From the IT Data Marts tree of SAS Data Integration Studio, double-click the job that contains the User-Written Staging transformation. The job opens on the **Diagram** tab of the Job Editor window.
2. Right-click the box that represents the User-Written Staging transformation.
3. From the drop-down menu, select **Properties**.
4. Click the tab that contains the information that you want to modify.
 - On the **General** tab, you can change the name of the transformation. The name of the transformation can be no more than 60 characters. It must be unique within the job that contains it.
 - On the **Staging Parameters** tab, you can change the standard staging parameters and the user-written staging parameters. For information about these parameters, see [“Create a User-Written Staging Transformation” on page 252](#).
5. The remaining tabs on the User-Written Staging Properties dialog box consist of standard SAS Data Integration Studio properties that can also be modified. For information about how to use these properties, click **Help** on the tab.
6. Click **OK** to return to the process flow diagram.
7. Click **File** ⇒ **Save** to save the job.

Rename a User-Written Staging Transformation

To rename a User-Written Staging transformation, perform the following steps:

1. From the IT Data Marts tree of SAS Data Integration Studio, double-click the job that contains the User-Written Staging transformation. The job opens on the **Diagram** tab of the Job Editor window.
2. Right-click the box that represents the User-Written Staging transformation.
3. From the drop-down menu, select **Properties**.
4. In the **Name** field on the **General** tab, enter the new name of the transformation. The name of the transformation can be no more than 60 characters. It must be unique within the job that contains it.

The new name of the transformation is displayed in the box that represents the User-Written Staging transformation in process flow diagram.

5. Click **OK** to return to the process flow diagram.
6. Click **File** ⇒ **Save** to save the job.

Using the Additional Generated Code

When the SAS code for the User Written Staging transformation is generated, additional code is generated and precedes the user-written code. This additional code provides information that can be useful in the user-written code. The generation of this additional

code is controlled by the options on the **User Written** section of the **Staging Parameters** tab.

Macro Variables

Macro variables for any extended attributes can be entered on the **Extended Attributes** tab of the User-Written Staging Properties dialog box. The name of each extension is the macro variable name, and the value of the extension is the value of the macro variable. If the name has spaces or other invalid characters, they are removed.

&input

If there is a source table to the transformation, then this macro variable contains the **libref.tablename** of the source table.

&target1-&targetn

These macro variables contain the names of the *n* target tables in the *libref.tablename* format.

Macros

These macros are available for use with user-written staging transformations.

- The following macro generates table options for specified target tables.

```
%macro _ITMS_tableOptions (targetTableNum=,
                           generateTableName=YES,
                           generateIndexes=YES);
```

- If the TARGETTABLENUM= parameter is left blank, then the table options for all the target tables are generated.
- If you specify **generateTableName=NO**, the macro generates a list of the table options.
- If you specify **generateTableName=YES**, the macro generates the table name with the table options in parentheses after it.
- The following macro generates an ATTRIB statement for each column of the requested table. If the TARGETTABLENUM= parameter is left blank, the macro generates an ATTRIB statement for each column for all tables. The ATTRIB statement includes the column name, type, length, format, informat, and label.

```
%macro _ITMS_attrib (targetTableNum=);
```

- The following macro generates assignment statements for all the columns of the requested table that will set the value to " " or . (missing), depending on the type of column. If the TARGETTABLENUM= parameter is left blank, the macro generates assignment statements for all columns in all tables.

```
%macro _ITMS_assignToMissing (targetTableNum=);
```

- The following macro generates a statement for each computed column of the requested table. If the TARGETTABLENUM= parameter is left blank, the macro generates a statement for each computed column for all tables.

```
%macro _ITMS_computedColumns (targetTableNum=);
```


- The following macro generates a list of all the columns of the requested table. If the TARGETTABLENUM= parameter is left blank, the macro generates a list of all the columns for all tables.

```
%macro _ITMS_columnList (targetTableNum=,
                        delimiter=,
                        sqlAlias=);
```

The additional parameters allow the list to be generated in different ways:

DELIMITER=

This parameter is set to a character string that will be put between each column in the list. The most common value of this parameter is a comma, so that the list can be used in SQL code. If a value is not specified, then the column names will be separated by blanks.

SQLALIAS=

This parameter is an alias that can be used as a prefix for each column name. There will be a period between this prefix and the column name. The most common use of this parameter is in SQL when a table alias is needed for each column.

Tabinfo and Varinfo

You can choose to create the Tabinfo and Varinfo tables in the Work library. These tables can be used in the user-written code to generate SAS code that is based on the target table and column information.

The Tabinfo table contains one record for each target table. The following columns are included in the Tabinfo table:

TABLENM

specifies the name of the table.

LIBREF

specifies the libref for this table.

ISCOMPRESSED

specifies the COMPRESS data set option. The values can be YES (that is, COMPRESS=YES) or NO (that is COMPRESS=NO).

INDEXOPTION

specifies the INDEX data set option (if there is an index on the table).

The Varinfo table contains one record for each column in all the target tables. The following columns are included in the Varinfo table:

TABLENM

specifies the name of the table.

VARNM

specifies the name of the column.

VARLAB40

specifies the column label.

VARFMT

specifies the column format.

VARINFMT

specifies the column informat.

VARLEN

specifies the length of the column.

VARTYPE

specifies the type of the column. The type can be C (that is, character) or N (that is, numeric).

FORMULA

specifies the expression that is used to create this column, if it is a computed column.

Chapter 12

Jobs

Jobs That Process IT Data	263
Overview of Jobs That Prepare, Stage, Aggregate, Generate Information Maps, and Report for IT Data	263
Processing IT Data Using the Adapter Setup Wizard	264
Processing IT Data without the Adapter Setup Wizard	267
Working with Jobs	268
Create a Job	268
Delete a Job	269
Modifying a Job	269
Rename a Job	273
Running Jobs	273
Run a Job Immediately	273
Deploy a Job for Scheduling	274
Running an “Overall” Job with the %RMRUNETL Macro	282
Allocating the Raw Data File Externally	282

Jobs That Process IT Data

Overview of Jobs That Prepare, Stage, Aggregate, Generate Information Maps, and Report for IT Data

To prepare and process IT resource data, SAS IT Resource Management requires jobs that perform the following tasks:

- Prepare the raw data for staging
- Stage the raw data
- Aggregate the data
- Generate information maps of the data
- Generate reports

For information about how to generate the job that prepares the raw data for staging, see [“Data Sources Supported by SAS IT Resource Management Adapters”](#) on page 290. The remaining jobs (staging, aggregating, generating information maps, and reporting) can be generated in two ways:

- automatically, by using the Adapter Setup wizard

- manually, by using the ITRM transformations

Note: After the jobs are created, either automatically by using the Adapter Setup wizard or manually by using the ITRM transformations, only the metadata definitions of the tables exist. You must actually execute the jobs in order to load data into the physical tables.

CAUTION:

Make sure that all tables that are associated with the job are closed before you run the job. If any one table cannot be accessed with Write access, the job will fail with a nonzero return code.

Processing IT Data Using the Adapter Setup Wizard

Generating Jobs with the Adapter Setup Wizard

You can automatically generate the metadata for the jobs that process IT resource data by invoking the Adapter Setup wizard.

To invoke the Adapter Setup wizard, from the SAS IT Resource Management menu select **File** ⇒ **New** ⇒ **IT Resource Management** ⇒ **Adapter Setup**. The Adapter Setup wizard opens and prompts you to specify information about the adapters that you want to work with. For any adapter, you can specify one of the following configuration levels that governs the types and number of jobs that the wizard generates.

- **Staging Job Only** creates only the staging job for the selected adapter.
- **Minimal** creates the staging job and only a small number of aggregations for the metrics that are used to generate a minimal set of reports for the selected adapter.
- **Typical** creates the staging job, all output from the Minimal level, and any additional metrics, aggregations, and report definitions that the SAS IT Resource Management solution has predetermined to be of interest for the selected adapter and reports. This option is the default setting.
- **Full** creates the staging job, all output from the Typical level, and any additional metrics, aggregations, and report definitions that make up a comprehensive set of jobs that are supported by the Adapter Setup wizard for the selected adapter.

When you click **Finish**, the Adapter Setup wizard creates the required jobs and all the metadata that is necessary for those jobs to execute. Among other information that it requests, the Adapter Setup wizard requires you to specify the IT data mart where it stores (by default) the jobs and the associated metadata for the adapter that you are working with. These jobs and the associated metadata are stored in the **<adapter><number>** folder that you are working with.

The Adapter Setup wizard creates the following jobs:

- one staging job that contains only the part of the ETL flow that relates to the staging transformation and the target staged tables.
- zero or more aggregation jobs that each contain one Aggregation transformation with target summary tables. The number of aggregation jobs that are created varies based on the domain categories (and their respective subcategories), time periods, and reporting selections that you made when specifying the adapter.
- zero or more information map jobs that each contain one Information Map transformation per aggregation table. The number of information map jobs that are created varies based on the domain categories (and their respective subcategories), time periods, and reporting selections that you made when specifying the adapter.

Note: In SAS IT Resource Management 3.1 and SAS IT Resource Management 3.1.1, the Adapter Setup wizard created jobs that created or updated aggregation tables and also generated information maps based on those tables. If the job that modified aggregation tables was run, the corresponding Information Map transformations were part of that job. Therefore, the information maps were automatically regenerated. Starting with SAS IT Resource Management 3.2, jobs that create or update aggregation tables do not include transformations that generate information maps. Therefore, if any changes are made to the tables that are input to those Information Map transformations that alter the structure of the table (for example, by adding or deleting columns), then the user must rerun the corresponding information map jobs so that the information maps are regenerated.

- zero or more reporting jobs that each contain one or more Gallery transformations per information map. The number of reporting jobs that are created varies based on the domain categories (and their respective subcategories), time periods, and reporting selections that you made when specifying the adapter.

Note: The jobs that the Adapter Setup wizard creates depend on one another to process the raw data for reporting. For example, the aggregation jobs include aggregations that rely on the staged tables that are generated when the staging job is run. As a result, each job should be run only after the jobs that it depends on have been run. The following list shows the types of jobs that the Adapter Setup wizard creates and the dependencies of each. Each job type in the list is dependent on the job type that precedes it.

1. staging job
2. aggregation job
3. information map job
4. reporting job

For information about how to work with the Adapter Setup wizard, see [Chapter 9, “Adapter Setup Wizard,”](#) on page 177.

Where Does the Adapter Setup Wizard Store the Objects That It Generates?

The Adapter Setup wizard stores the objects that it generates in predetermined locations within the IT data mart that you select. The wizard generates two sets of subfolders for the objects that it creates:

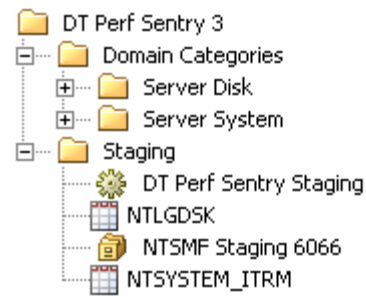
- the **Staging** subfolder, which contains the staging job, staged tables, and the library for those tables
- the **<Domain Category>** subfolder within the **Domain Categories** folder, which contains the aggregation, information map, and reporting jobs, their associated tables and maps, and the library that is associated with those objects

However, jobs, tables, and libraries can be stored in any location to which you have write access. To change where an object is stored, right-click the object (in the process flow diagram or in the tree where it is currently stored) and select **Properties**. On the **General** tab of the Properties dialog box, click the **Browse** button that is next to the **Location** field. The Select a Location dialog box appears. Navigate to the location that you want to use and click **OK**.

The default location for the staging job, the staged tables, and the associated library is the **Staging** subfolder that is located within the adapter folder of the IT data mart that you selected.

The following display shows the metadata that is created in the **Staging** subfolder for the third instance of the DT Performance Sentry adapter within the IT data mart. The **Staging** subfolder contains the DT Perf Sentry Staging job, the NTLGDSK and NTSYSTEM_ITRM staged tables, and the library for those tables—NTSMF Staging 6066.

Display 12.1 Default Location of Objects Created by the Adapter Setup Wizard for Staging the Raw Data




The default location for the remaining jobs, their associated tables, and library is in the **Domain Categories** subfolder that is located within the adapter folder of the IT data mart that you selected.

The following display shows the metadata that is created in the **Domain Categories** subfolder for the third instance of the DT Performance Sentry adapter within the IT data mart. The **Domain Categories** subfolder contains two subfolders of domain objects: Server Disk and Server System. Each domain subfolder contains five aggregated tables,

the library associated with those tables, the aggregation job, the information map job, and the two reporting jobs.

Display 12.2 Default Location of Objects Created by the Adapter Setup Wizard for Aggregating, Creating Information Maps, and Reporting on the Staged Data



Note: Information maps are not generated until the information map job runs. When they are generated, they are stored by default in the appropriate **Domain Categories** subfolder, although you can store them in any location to which you have Write access.

Processing IT Data without the Adapter Setup Wizard

You can manually generate the metadata for the tasks that process IT resource data instead of invoking the Adapter Setup wizard. To do so, you must first open a job. Then, within that job, you can specify the staging, aggregating, information map, and reporting transformations that are needed. You can also modify existing jobs by adding and deleting transformations, tables, and information maps.

To generate the metadata for a transformation, perform the following steps:

1. Open a job. If you want to work with an existing job, double-click on that job to open it on the **Diagram** tab of the Job Editor window of SAS IT Resource Management.

If you need to create a job, you can invoke the New Object wizard from the **File** menu of SAS IT Resource Management. For information about how to perform this task, see [“Create a Job” on page 268](#).

The new job opens on the **Diagram** tab of the Job Editor window.

2. Navigate to the **ITRM Transformations** folder of the Transformations tree of SAS IT Resource Management.
Note: A job can consist of a single transformation, or it can include multiple transformations such as staging transformations, Aggregation transformations, Information Map transformations, and Gallery transformations.
3. Drag the transformation that you want to work with onto the job in the **Diagram** tab of the Job Editor window.
4. Locate and follow the instructions for working with transformations that perform the following tasks:
 - Stage the raw data
 For information about how to stage the raw data, see [Chapter 7, “Staging the Data,” on page 105](#).
 - Aggregate the data
 For information about how to aggregate the data, see [Chapter 8, “Aggregating the Data,” on page 131](#).
 - Generate information maps of the data
 For information about how to generate information maps, see [on page 207](#).
 - Report on the IT intelligence and analysis that SAS IT Resource Management generates
 For information about how to generate, manage, and view reports, see the *SAS IT Resource Management : Reporting Guide*.

Note: In order to create (or update) the tables that are specified by the transformation, the job that contains the transformation must be executed. Jobs can be run immediately from SAS IT Resource Management, or they can be scheduled to execute as part of the operational procedures of an organization. (Due to the size of the aggregation jobs, it is preferable to run them in batch mode.)

Working with Jobs

Create a Job

To create a job, perform the following steps:

1. From the SAS Data Integration menu bar, select **File** ⇒ **New** ⇒ **Job** to open the New Job wizard.
2. In **Name** field, you can override the default name that is assigned to the job. You can enter up to 60 characters in this field. A value is required in this field.
3. In **Description** field, enter the information that describes the job. You can enter up to 200 characters in this field. A value is optional in this field.
4. In **Location** field, you can override the default location of the job. To do so, click **Browse** to open the Select a Location dialog box where you can navigate to another location.

Jobs must be created from within a folder in an IT data mart. If the New Job wizard is not invoked from within a folder in an IT data mart, SAS IT Resource

Management creates a default folder (**NewFolder**) within the last referenced IT data mart. For information about how to create an IT data mart, see “[Create an IT Data Mart](#)” on page 57. To create a folder within an IT data mart, perform the following steps:

Create a Folder

- a. Navigate to the IT data mart.
 - b. From the SAS Data Integration menu, select **File** ⇒ **New** ⇒ **Folder**. The New Folder wizard opens.
 - c. (Optional) Override the default name of the folder with a name that does not exceed 60 characters. A value is required in this field.
Note: For best results, enter a name that indicates the purpose of the folder.
 - d. (Optional) Enter a description of the folder. You can enter up to 200 characters. A value is optional in this field.
 - e. (Optional) Override the default location by clicking **Browse** to open the Select a Location dialog box, where you can navigate to another location. A value is required in this field.
 - f. Click **OK** to create the folder and store it in the IT data mart from which you invoked the New Folder wizard.
5. Click **OK** to create the job and store it in the location that you chose.

Note: The **Diagram** tab of the Job Editor window of SAS IT Resource Management automatically displays the job that you created. (The window is empty except for the job name that displays on the banner of the window.)

6. To view the new job in the folder, click **Refresh** in the menu bar. (In order to activate the **Refresh** icon in the menu bar, click anywhere in the left panel of SAS IT Resource Management.)

Delete a Job

To delete a job, perform the following steps:

1. Navigate to the location of the job in the IT Data Marts tree.
2. Right-click the job.
3. From the menu that is displayed, select **Delete**.
4. In the confirmation box, click **Yes** to delete the job.

Modifying a Job

Add a Transformation to a Job

To add a transformation to a job, perform the following steps:

1. Navigate to the location of the job in the IT Data Marts tree.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor window.
3. In the Transformations tree, expand the **ITRM Transformations** folder.

4. Drag and drop the transformation that you want to add to the job onto the process flow diagram.
5. Right-click the transformation to display a menu of options with which you can specify the transformation that you are adding to the job.
 - For information about how to add a staging transformation, see [“Working with Staging Transformations” on page 107](#).
 - For information about how to add an Aggregation transformation, see [“Working with Aggregation Tables” on page 140](#).
 - For information about how to add an Information Map transformation, see [“User-Specified Information Map Transformations” on page 212](#).
 - For information about how to add a Gallery transformation, see the “Working with the Gallery Transformation” topic in the “Running Reports” chapter of the *SAS IT Resource Management: Reporting Guide*.
6. Save the job that you modified. To do so, select **File** ⇒ **Save**.

Delete a Transformation from a Job

To delete a transformation from a job, perform the following steps:

1. Navigate to the location of the job in the IT Data Marts tree.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor.
3. Right-click the transformation that you want to delete.
4. Save the job that you modified. To do so, select **File** ⇒ **Save**.

Add a Table to a Transformation

You can add a table to a staging transformation or an Aggregation transformation. To add a table to a transformation in a job, perform the following steps:

1. Navigate to the location of the job in the IT Data Marts tree.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor.
3. Right-click the transformation to which you want to add a target table. A menu of options is displayed.
 - The menu for a staging transformation provides an action to **Add New Staged Table**. For information about this topic, see [“Working with Staging Transformations” on page 107](#).
 - The menu for an Aggregation transformation provides an action to **Add Aggregation Table**. For information about this topic, see [“Working with Aggregation Tables” on page 140](#).
4. Select the table that you want to add to invoke the wizard that will generate the metadata for that table.
5. The wizard prompts you for information that is required to specify the metadata for the table. When you are satisfied with your specifications, click **Finish** to generate the metadata for the table.
6. Save the job that you modified. To do so, select **File** ⇒ **Save**.

Delete a Table from a Transformation in the Process Flow Diagram

To delete a table from a transformation in the process flow diagram, perform the following steps:

1. Navigate to the location of the job in the IT Data Marts tree.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor.
3. Right-click the table that you want to delete. The table is removed from the transformation. However, the metadata of this table and its library are not deleted from the IT Data Marts tree, and its contents are not deleted from the physical file.
4. Save the job that you modified. To do so, select **File** ⇒ **Save**.

If you delete the source table of an Aggregation transformation in the process flow diagram, the target tables of that transformation are also deleted. However, if this action is not acceptable, you can reinstate these tables. To do so, from the **Edit** menu, click **Undo**.

CAUTION:

To avoid processing errors, delete any reference to a deleted table, as a source or target for any transformation in any other job. For example, Information Map transformations that use the deleted table as a source table should be deleted.

Delete a Table from the IT Data Marts Tree

The **Delete** action deletes the metadata for the table in the IT Data Marts tree. The physical table and the library are not deleted. To delete a table from the IT Data Marts tree, perform the following steps:

1. Navigate to the folder that contains the table in the IT Data Marts tree.
2. Expand that folder and right-click the table that you want to delete.
3. From the menu that appears, select **Delete**.
4. In the confirmation box, click **Yes**.

CAUTION:

A job that accesses a table that is deleted will not run successfully.

Delete a Table with Its Contents from a Transformation in the Process Flow Diagram

The **Delete with Contents** action deletes the metadata of the table and the physical table. It does not delete the associated library. To delete a table with its contents from the IT Data Marts tree, perform the following steps:

1. Navigate to the folder that contains the table in the IT Data Marts tree.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor.
3. Right-click the table whose contents and metadata you want to delete.
4. From the menu that displays, select **Delete with Contents**.

CAUTION:

A job that accesses a table that is deleted will not run successfully.

Delete a Table with Its Contents from the IT Data Marts Tree

The **Delete with Contents** action deletes the metadata of the table and the physical table. It does not delete the associated library. To delete a table with its contents from the IT Data Marts tree, perform the following steps:

1. Navigate to the folder that contains the table in the IT Data Marts tree.

2. Expand that folder and right-click the table whose contents and metadata you want to delete.
3. From the menu that displays, select **Delete with Contents**.

CAUTION:

A job that accesses a table that is deleted will not run successfully.

Erase a Table from a Transformation

Erasing a table deletes the physical table, its contents, and its metadata. To erase a table, perform the following steps:

1. In a tree view of SAS IT Resource Management, locate the table that you want to erase.
2. Right-click the table that you want to erase. The job that uses the table must be closed and currently not in use before erasing.
3. Select **Erase**.
4. Click **Yes** to confirm that you want to erase this table.

CAUTION:

If you erase a table that serves as a source table in other jobs, then these associated jobs might not execute because they are missing a source table.

Simply erasing a table does not affect the metadata of the associated jobs. However, if you erase a table and then rerun the job that creates that table, then the data for the deleted table is not created and the jobs that used that table will fail due to a missing source table.

Purge a Table from a Transformation in the Process Flow Diagram

The **Purge** action removes all observations from physical tables but leaves the physical table structure intact. To purge a table from the process flow diagram, perform the following steps:

1. Navigate to the job that contains the table in the IT Data Marts tree.
2. Double-click the job to open it on the **Diagram** tab of the Job Editor.
3. Right-click the table whose contents you want to purge.
4. From the menu that displays, select **Purge**.
5. In the confirmation box, click **Yes**.

Purge a Table from the IT Data Marts Tree

The **Purge** action removes all observations from physical tables but leaves the physical table structure intact. To purge a table from the IT Data Marts tree, perform the following steps:

1. Navigate to the folder that contains the table in the IT Data Marts tree.
2. Expand that folder and right-click the table that you want to purge.
3. From the menu that displays, select **Purge**.
4. In the confirmation box, click **Yes**.

Rename a Job

To rename a job, perform the following steps:

1. Navigate to the folder that contains the job in the IT Data Marts tree.
2. Right-click the job.
3. From the menu that displays, select **Rename**.
4. In the highlighted field, enter the new name of the job.

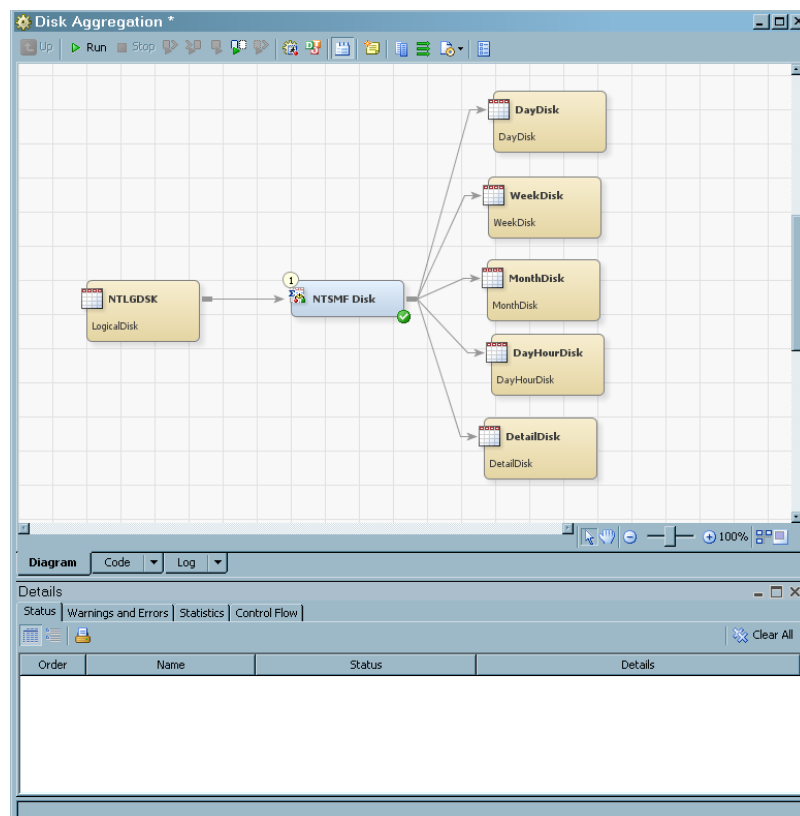
Running Jobs

Run a Job Immediately

After a job is created, and the transformation within it is defined, the job is ready to be run. To run a job immediately on the SAS application server where you are currently logged on, perform the following steps:

1. If the job is not open on the **Diagram** tab of the Job Editor window, open it by navigating to the job in the IT Data Marts tree and double-clicking it.
2. In the toolbar of the Job Editor window, select **Run**. This action submits the job's source code.

Display 12.3 Process Flow of an Aggregation Job in the Job Editor Window



3. The **Diagram**, **Code**, and **Log** tabs are located at the bottom of the window. You can select the **Code** tab to view the generated code. After the job executes, you can select the **Log** tab to view the SAS log.

Deploy a Job for Scheduling

How to Deploy a Job

You can deploy a job so that it can be scheduled to run in batch mode on the default SAS application server.

To deploy a job to be run in batch mode, perform the following steps:

1. In the IT Data Marts tree of SAS IT Resource Management, navigate to the job that you want to deploy.

Note: By pressing the CNTL key, you can select multiple jobs within an IT data mart to deploy. However, if the server is z/OS, you can deploy only one job at a time.

2. Right-click the job and from the menu that appears, select **Scheduling** ⇌ **Deploy**. The Deploy a job for scheduling dialog box appears.

- Under Windows or UNIX operating environments, the following dialog box appears.

Display 12.4 Deploying a Windows or UNIX Job for Scheduling

The screenshot shows a Windows-style dialog box titled "Deploy a job for scheduling". It contains the following elements:

- Batch Server:** A dropdown menu currently set to "SASITRM - SAS DATA Step Batch Server".
- Deployment Directory:** A dropdown menu currently set to "Batch Jobs", with a "New..." button to its right.
- Deployed Job Name:** A text input field containing the text "Server_Disk_Aggregation".
- Location:** A text input field containing a long file path: "/Shared Data/SAS IT Resource Management/IT Data Marts/My IT Data Mart/DT Perf Sentry 1/Domain Categories/Server Disk". To its right is a "Browse..." button.
- Buttons:** At the bottom right, there are three buttons: "OK", "Cancel", and "Help".

- Under the z/OS operating environment, the following dialog box appears.

Display 12.5 Deploying a z/OS Job for Scheduling

Deploy a job for scheduling

Batch Server: SASITRM-SAS DATA Step Batch Server

Deployment Directory: Batch Jobs New...

Deployed Job Name: System_Aggregation

Location: /Shared Data/SAS IT Resource Management/IT Data Marts/AAA/SAR 1/Domain Categories/System Browse...

JCL

Deployment Directory: Batch Jobs

```
// $JOBNAME$ JOB ACCOUNT-CODE, 'PROGRAMMER-NAME'
//      MSGCLASS=A, MSGLEVEL=(1,1),
//      TIME=(2,00), REGION=64M
// *****
// *   SAS SCHEDULING TEMPLATE JCL
// *****
// *
// *   SAS LIBRARY CONTAINING PROCS
// *
//SASLIB  JCLLIB  ORDER=CPD.ITRM32.DJ.V92D8022.CNTL
// *
// * DO NOT MODIFY THE FOLLOWING LOGSTART STEP.
// * IT IS NECESSARY FOR PROPER JOB SCHEDULING.
// *
// * $LOGSTART EXEC PGM=BPXBATCH, PARM='SH'
// * $STRTM DD PATHNAME
```

OK Cancel Help

3. The batch server, the deployment directory, and the name and path of the job that you selected are the default values that are displayed in this window.

- The **Batch Server** field contains the server where the job is to be executed.
- The **Deployment Directory** field specifies the physical path to the directory where the generated code for the job is stored. From the drop-down list, select a different deployment directory. Select **New** to open the New directory window where you can define a new deployment directory.
 - If you define a new deployment directory in the Windows or UNIX operating environments, the new deployment directory will contain the SAS source code for the job.
 - If you define a new deployment directory in the z/OS operating environment, the new directory will contain a file that has the SAS source code for the job and a file that contains the generalized job control language (JCL) that can run the job. You can modify this generalized file of JCL according to the requirements of your site.

Note: At this time, the JCL Deployment Directory must be deployed to a UFS location and not a traditional file system (PDS) because partitioned data sets are not yet supported for JCL Deployment directory.


- The **Deployed Job Name** field specifies the name of the file that contains the generated code for the selected job. You can edit the default name, but the name must be unique in the context of the Deployment Directory that was previously specified. This field is available only when you deploy a single job. When you deploy multiple jobs, the filename is the same as the job names, but the name has a .sas extension.

Because the default value of the **Deployed Job Name** that is displayed is based on the name of your Job object, you might need to revise the name in order to make it a legal filename on the operating system to which it is being deployed.

Note: For z/OS traditional file system locations, if the **Deployment Directory** name points to a partitioned data set (PDS), then the deployed job name specifies a member name for that PDS. The member name of the PDS cannot exceed 8 characters.

- The **Location** field contains the path to the Folders tree that contains the job. Click **Browse** to open the Select a location dialog box where you can select another location.

Note: For more information about the fields on this window, click **Help**.

4. Click **OK** to deploy the job.
5. The code is then generated for the job and stored in the directory that is specified. If the deployment is successful, a message is displayed confirming that result. Metadata about the deployed jobs is saved to the current SAS Metadata Server. A status window is displayed and indicates whether the deployment was successful. Another object with the same name as the deployed job is stored in the IT Data Marts tree. This deployed job can be identified by the deployed job icon: . This deployed job is now available for scheduling.

If you update a job or if your computing environment changes, you must redeploy it so that the latest version of the job is the available for execution. For example, if you add, remove, or change a transformation in a job, you must redeploy it to have those changes take effect.

Schedule a Job to Run in Batch

After you deploy a job for scheduling, you can deploy it using any scheduling software that is available in your operating environment. An administrator can set up and run your job with any of the following types of software:

- any scheduler that is built in to your operating system.
- any third-party scheduler.
- the Schedule Manager plug-in to SAS Management Console. This software can define a flow, add one or more deployed jobs to it, and schedule the flow to run.

Note: For more information about scheduling, see *Scheduling in SAS*. This document is associated with plug-ins to SAS Management Console. It is available at <http://support.sas.com/documentation/onlinedoc/sasmc/index.html>.

How to Schedule a Job Using a Built-In or Third-Party Scheduler

The deployed job is a file that contains SAS code. In order to schedule it using a built-in or third-party scheduler, you must supply an appropriate invocation of SAS, which runs this code.

- On UNIX, you should supply a shell script.
- On Windows, you should supply a batch file.
- On z/OS, you should supply the JCL. A file of JCL that can be modified according to the requirements of your site can be found in the deployment directory that you selected.

This invocation must invoke the correct version of SAS with the configuration that is specified in the SAS IT Resource Management installation and configuration instructions. For more information about how to invoke SAS on your operating system, see *SAS Companion for UNIX Environments*, *SAS Companion for z/OS*, or *SAS Companion for the Microsoft Windows Environment*, as appropriate.

After you have set up the SAS invocation, you can then use your preferred scheduler to add this job to the script, batch file, or JCL that invokes SAS. For example, on UNIX you can use 'cron' and add the shell script invocation to the **crontab** file. For third-party schedulers and for schedulers that are supplied as part of the operating system, follow the documentation for that scheduler.

How to Schedule a Job Using the Schedule Manager in SAS Management Console

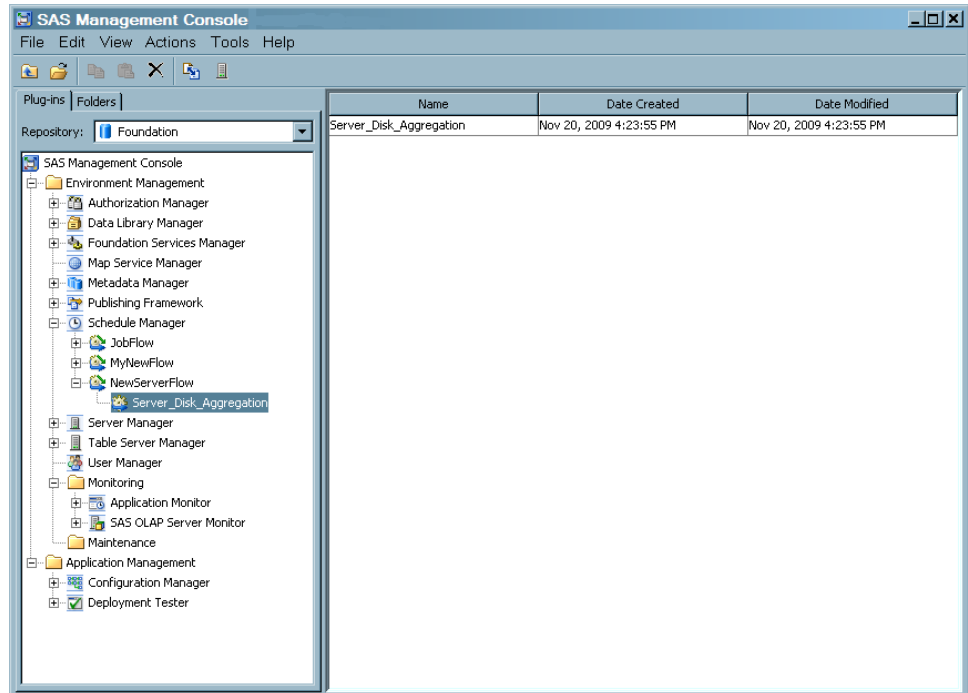
To schedule a job to run in batch mode, perform the following steps:

1. On the main navigation tree of SAS Management Console, right-click **Schedule Manager**. From the drop-down list, select **New Flow** to open the New Flow window.


Display 12.6 New Flow Window in SAS Management Console

The screenshot shows the 'New Flow' dialog box. The 'Name' field is set to 'NewServerFlow'. The 'Location' is 'Shared Data'. The 'Scheduling Server' is 'Operating System Services'. The 'Available Items' list contains 'DailyJob' and 'MonthlyJob'. The 'Selected Items' list contains 'Server_Disk_Aggregation'. The 'View Properties' button is at the bottom left. The 'OK', 'Cancel', and 'Help' buttons are at the bottom right.


2. In the **Name** field, enter the name of the job flow. (The name cannot contain spaces.)
3. If you want to specify a new scheduling server, select **New**. The New Scheduling Server wizard opens.
4. From the **Available Jobs** column, select the jobs that you want to schedule. Use the arrow to transfer those jobs to the **Selected Jobs** column. (In the preceding display, the NewServerFlow called Server_Disk_Aggregation job is selected for scheduling.)
5. Click **Properties** to specify the scheduling details and authorization requirements of the jobs.
6. Click **OK** to define the job flow.

Display 12.7 SAS Management Console Window Showing Scheduled Job Flow

7. Select the new flow under **Schedule Manager**.
8. (Optional) Highlight the flow and from the **Actions** menu, select **Manage Dependencies**. The Manage Dependencies dialog box appears, where you can specify one or more conditions that control the starting of the current job.

Display 12.8 Manage Dependencies Dialog Box

9. From the **Actions** menu, select **Schedule Flow**. The Schedule Flow dialog box appears.

Display 12.9 Select a Trigger to Start the Current Flow

10. From the drop-down list, select when to run the flow—to run now or to run manually. Click **Options** to open the Trigger Options dialog box on which you can select a trigger to start the flow, such as a specific date, time, and recurrence. Click **OK**. When prompted, enter a user ID and password for the scheduling server. A message will be displayed confirming the successful scheduling of the job.

For information about the functions that can be performed from SAS Management Console, and options that can be specified when scheduling flows, see the online Help for the Schedule Manager plug-in to SAS Management Console and the *SAS Management Console: User's Guide*.

Redeploy a Job

If you update a job that you already deployed, you must redeploy it so that the latest version of the job is scheduled. For example, if a table is deleted from a job, then the job must be deployed again.

To redeploy a job, perform the following steps:

1. In the IT Data Marts tree of SAS IT Resource Management, navigate to the job that you want to redeploy. You can select multiple jobs within an IT data mart to redeploy.
2. Right-click the job and from the menu that displays, select **Scheduling** ⇒ **Redeploy**.

The Redeploy a job for scheduling dialog box appears.

- Under Windows or UNIX operating environments, the following dialog box appears.

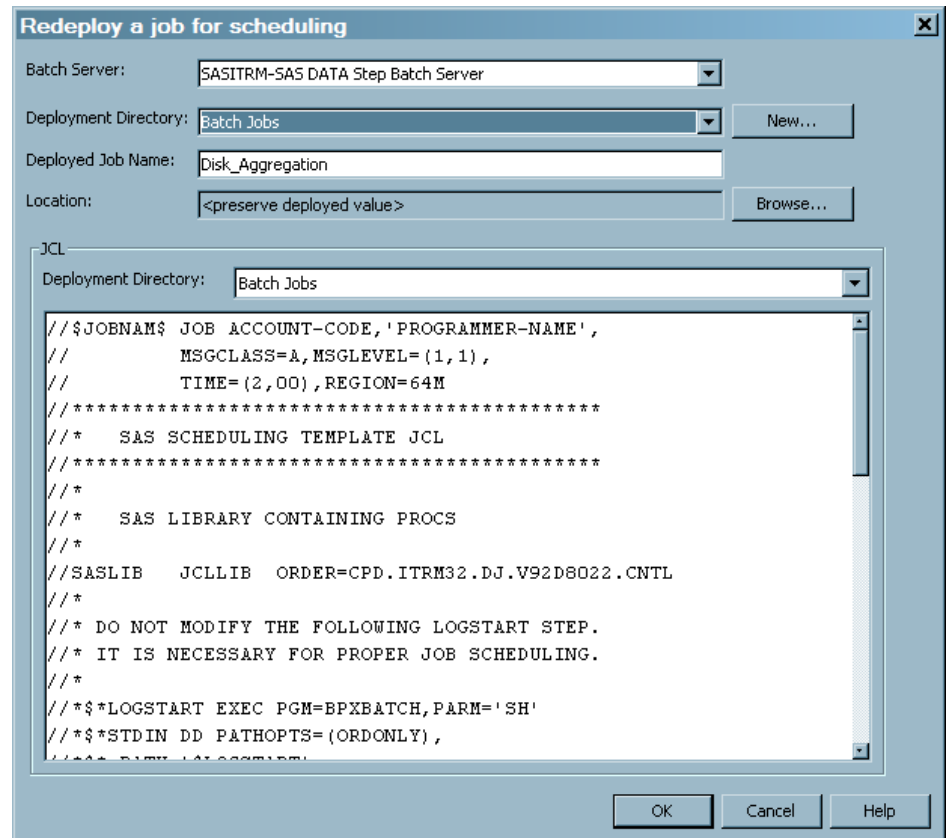
Display 12.10 Redeploying a Windows or UNIX Job

You can choose to preserve the deployed value of the job's metadata. However, you can change where the job is redeployed by changing the deployment directory and its location. If you define a new deployment directory in the Windows or UNIX operating environments, the new deployment directory will contain the SAS source code for the job.

A message displays the name of the jobs that will be redeployed. These jobs are all the deployed jobs that are contained on the server that you are connected to. Click **Yes** to redeploy them.

- Under the z/OS operating environment, you can also choose to preserve the deployed value of the job's metadata. If you want to change where the job is redeployed, select a new value for the **Deployment Directory** field. If you define a new deployment directory in the z/OS operating environment, the new directory will contain the SAS source code for the job and a file that contains the generalized JCL that can run the job. You can modify this generalized file of JCL according to the requirements of your site.

The following dialog box appears.

Display 12.11 Redeploying a z/OS Job

3. A message appears that lists the jobs that were redeployed. If a job was not redeployed, an error message appears.

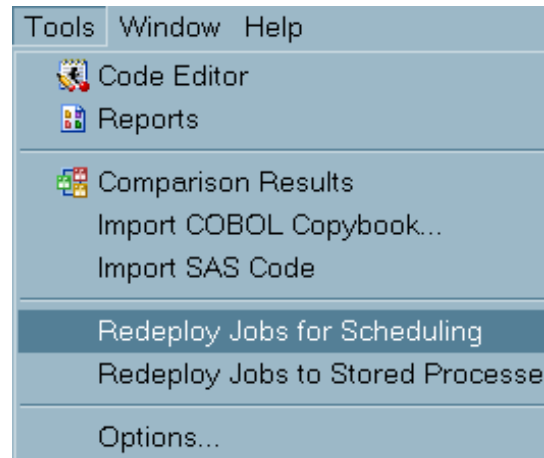
Redeploy All Jobs on the Server

If you made changes that affect all jobs that are deployed on the server, you must redeploy then so that the latest versions of the jobs are scheduled. For example, if the computing environment of the server where the deployed jobs are stored changes, the jobs must be deployed again.

Note: The code that is generated for a job contains the credentials of the person who created the job. If those credentials have changed, and a deployed job contains outdated user credentials, the deployed job will fail to execute. In that case, redeploy the job with the appropriate credentials.

To redeploy jobs, perform the following steps:

1. From the SAS IT Resource Management menu bar, select **Tools** ⇒ **Redeploy Jobs for Scheduling**.

Display 12.12 Tools Menu of SAS IT Resource Management

A message then displays the names of the jobs that will be redeployed. These jobs are all the deployed jobs that are contained on the server that you are connected to. Click **Yes** to redeploy them.

2. Under the z/OS operating environment, you can choose to preserve the deployed value of the job's metadata. If you want to change where the job is redeployed, select a new value for the **Deployment Directory** field. If you define a new deployment directory in the z/OS operating environment, the new directory will contain the SAS source code for the job and a file that contains the generalized JCL that can run the job. You can modify this generalized file of JCL according to the requirements of your site. Code is generated for the jobs that you redeployed. The code is saved to the deployment directory of the SAS application server that is used to deploy jobs.

Running an “Overall” Job with the %RMRUNETL Macro

You can choose to run your staging, aggregation, information map, and reporting jobs within a single processing macro. The staging job is run first, and then the aggregation jobs are run in parallel using Multi-Processing (MP) Connect, followed by the information map jobs that are run in parallel, and finally the report jobs that are run in parallel. If the staging job fails, the aggregation, information map, and report jobs are not run. After each job completes, the SAS log is examined by the program. The completion status of the job (Success, Warning, or Error) is sent by e-mail to the specified recipient. If the log displays warnings or errors, the e-mail includes the text of the error or warning along with the corresponding line number in the log file.

For information about how to set up the %RMRUNETL macro that provides this ability, see “%RMRUNETL” on page 529.

Allocating the Raw Data File Externally

When a job is deployed, the file that was entered in the **Raw data file or directory** field of the **Staging Parameters** tab of the staging transformation (or in the Adapter Setup wizard) appears in the generated code as part of a macro assignment (a %let statement). In a production environment, it is possible that the raw data file changes from day to day, or that it needs to be allocated externally, for example in the JCL of a z/OS batch job.

Note: For adapters that accept files or directories as input, the code is %LET
RAWDATA=raw-data-assignment.

If the raw data file or directory is allocated externally, it will be used in preference over the location referred to in the `%let` statement. For completeness, the `%let` statement can be removed, but this is not necessary.

CAUTION:

Do not remove or comment out the `%let` statement for adapters that do not have a raw data file or a directory. For some adapters such as the SAP ERP, MS SCOM, HP Reporter, VMware Data Acquisition, and VMware vCenter adapters, the `%let` statement is present in the code. However, that statement does not specify a file or a directory. Instead, the `%let` statement points to the libref of the SAS library that is associated with the raw data. (The location of the input library typically remains fixed, unlike the name or location of individual raw data files.)

If the job is redeployed, the raw data assignment will be reset and will need to be removed again, unless the raw data parameter in the **Staging Parameters** tab is removed beforehand.

When allocating the raw data file externally for MXG based adapters, the correct fileref must be used. (On z/OS, the fileref is assigned with a DD JCL statement.)

For adapters that are not based on MXG, use a fileref of RAWDATA. For adapters that are based on MXG, use the filerefs in the following table:

Table 12.1 Table of Filerefs for Adapters That Are Based on MXG

Adapter	Fileref
IBM DCOLLECT	DCOLLECT
IBM EREP	EREP
BMC Mainview IMS	IMSLOG
DT Perf Sentry with MXG	NTSMF
IBM SMF	SMF
ASG TMONDB2	TMDBIN
ASG TMON2CIC	MONICICS
CA TMS	TMC
IBM TPF	TPFIN
IBM VMMON	MWINPUT

Part 3

Appendixes

<i>Appendix 1</i>	
SAS IT Resource Management Components	287
<i>Appendix 2</i>	
Data Sources Supported by SAS IT Resource Management Adapters	289
<i>Appendix 3</i>	
Staging Parameters	355
<i>Appendix 4</i>	
Jobs That the Adapter Setup Wizard Creates for Domain Categories	381
<i>Appendix 5</i>	
Duplicate-Data Checking	391
<i>Appendix 6</i>	
Adapter Key Metrics	417
<i>Appendix 7</i>	
Data Model	459
<i>Appendix 8</i>	
Naming Standards	503
<i>Appendix 9</i>	
Macros	509
<i>Appendix 10</i>	
Best Practices and Troubleshooting Tips	535
<i>Appendix 11</i>	
Statistics	553

Appendix 1

SAS IT Resource Management Components

Server Tier Components

SAS IT Resource Management Server for the AIX, HP PA-RISC, HP-UX on Itanium, Solaris SPARC and Solaris 10 x64, Linux 32-bit and x64, z/OS, and Windows x64 and 32-bit operating environments uses the following components:

- Base SAS
- SAS IT Resource Management Server
- MXG from Barry Merrill (z/OS only and included in the IT Resource Management license)
- Demand Technology Performance Sentry (Windows x64 and 32-bit)
- SAS Metadata Server
- SAS Data Integration Studio
- SAS Enterprise BI Server
- SAS/SHARE
- SAS/ETS
- SAS/STAT
- SAS HPF (includes SAS Threaded Kernel Extensions for Operations Research)
- SAS/ACCESS Interface to ORACLE
- SAS/ACCESS Interface to ODBC (for AIX, HP PA-RISC, HP-UX on Itanium, Linux 32-bit and x64, Solaris SPARC and Solaris 10x64, and Windows 32-bit and x64)
- SAS/ACCESS Interface to DB2 (for AIX, HP PA-RISC, HP-UX on Itanium, Linux 32-bit, Solaris SPARC, z/OS, Windows 32-bit, and Windows 64-bit for Itanium)
- SAS/ACCESS Interface to PC Files

Middle Tier Components

SAS IT Resource Management Web Tier for the AIX, HP-UX on Itanium, Solaris SPARC and Solaris 10 x64, Linux x64, z/OS, and Windows x64 and 32-bit operating environments uses the following components:

- SAS IT Resource Management Middle Tier
- SAS Web Report Studio

- SAS Information Delivery Portal Web Application
- SAS Web Parts for Microsoft SharePoint
- SAS BI Dashboard
- SAS BI Portlets (not available on z/OS)
- SAS Content Server

Note: An application server is obtained separately. Application servers that are supported include the following products:

- JBoss (AIX, HP-UX on Itanium, SPARC and Solaris 10 x64, Linux x64, and Windows x64 and 32-bit)
- Oracle WebLogic (HP UX on Itanium using a 32-bit JRE, Solaris SPARC and Solaris 10 x64 using a 32-bit JRE, and Windows 32-bit)
- IBM WebSphere (z/OS, AIX, Solaris SPARC and Solaris 10 x64, Linux x64, and Windows x64 and Windows 32-bit)

Client Tier Components

SAS IT Resource Management SAS Client for the Windows XP Professional, Windows Vista, and Windows 7 workstation operating environments uses the following components:

- SAS Data Integration Studio
- SAS IT Resource Management Client
- SAS Enterprise Guide
- SAS Management Console (AIX, HP PA-RISC, HP-UX on Itanium, Linux 32-bit and x64, Solaris SPARC and Solaris 10 x64, and Windows 32-bit and x64)
- SAS Add-In for Microsoft Office
- SAS Information Map Studio
- SAS Workflow Studio
- SAS OLAP Cube Studio
- SAS OLAP Server Monitor
- DataFlux Data Management Client

Appendix 2

Data Sources Supported by SAS IT Resource Management Adapters

Data Sources Supported by SAS IT Resource Management Adapters	290
Supported Adapters	291
Accessing Raw Data for MXG	298
Preparing MXG Data for SAS IT Resource Management	298
Extracting Data from MXG	298
Accessing Raw Data for HP Perf Agent	298
Preparing HP Perf Agent Data for Staging	298
Extracting Data from HP Perf Agent	298
Example 1	299
Example 2	300
Accessing Raw Data for DT Perf Sentry	300
Preparing DT Perf Sentry Data for SAS IT Resource Management	300
Preparing DT Perf Sentry Data for Processing on z/OS	301
Extracting Data from DT Perf Sentry	301
Notes about DT Perf Sentry	302
Accessing the ITRM Windows Server DCS	302
Accessing the ITRM Exchange DCS	303
Accessing Raw Data for BMC Perf Mgr	303
Preparing BMC Perf Mgr Data for SAS IT Resource Management	303
Extracting Data from BMC Perf Mgr	304
Notes about BMC Perf Mgr	305
Accessing Raw Data for SAR	306
Preparing SAR Data for SAS IT Resource Management	306
Extracting Data from SAR	306
Using the sar Command	307
Accessing Raw Data for SNMP	308
Sources of SNMP Data	308
Extracting Data by Using HP NNM Tools	310
Extracting Data by Using the Round-Robin Database Tool (RRDTOOL)	311
Notes about the SNMPWALK Command	312
Working with the ITMS_SNMP_DeviceType Lookup Table	313
Working with Management Information Base (MIB) Definitions for SNMP Staged Tables	314
Accessing Raw Data for Web Log	315
Preparing Web Log Data for SAS IT Resource Management	315
Accessing Raw Data for Database Adapters	315
What Are the Database Adapters?	315

Preparing a Database Adapter for Staging	315
Duplicate-Data Checking for Database Adapters	316
Notes about Accessing HP Reporter Data	317
Notes about Loading Tasks and Events Data for VMware	319
Working with the VMware Lookup Table	321
Accessing Raw Data for SAP ERP	323
Working with Custom Tables	328
IBM SMF Custom Tables	328
How to Load IBM SMF Data into User-Defined Tables	328
Working with MXG Code to Stage Data	330
About MXG Adapters	330
MXG Views	330
IMACWORK and RMFINTRV Customizations Relevant to the RMF Domain Category of the SMF Adapter	331
IMACSHFT Customizations Relevant to All Domain Categories of the SMF Adapter	331
Customizations Relevant to the Jobs Domain Category of the SMF Adapter . . .	331
Customizations Relevant to the CICS and DB2 Domain Categories of the SMF Adapter	332
Customizations Relevant to the Jobs Domain Category of the CA TMS Adapter .	333
Customizations Relevant to the ASG TMON2CIC Adapter	333
Specify Spin Libraries for ASG TMON2CIC and IBM SMF	334
Limit IBM SMF Records	336
How MSU and MIPS Columns Are Included in Staged Tables for the SMF Adapter	336
VMware vCenter Data Acquisition Jobs	351
Overview of VMware Data Acquisition Jobs	351
How Does the VMware Data Acquisition Job Work?	352
Preserving Information about the VMware Environment	352
When Should the VMware Data Acquisition Job Run?	353

Data Sources Supported by SAS IT Resource Management Adapters

SAS IT Resource Management supports various adapters that enable you to load raw data about IT resources from many diverse data sources and convert that raw data into a standard form that facilitates aggregation and reporting on your IT resources. These adapters extract data from raw data files in a given data source and load that data into staged tables in the IT data mart.

The following table lists each SAS IT Resource Management adapter and identifies the software versions of the raw data sources that were used to verify each of them. The remainder of this appendix provides information about how to prepare the raw data source for use by SAS IT Resource Management.

Note: The specific software versions of the raw data sources in this table have been used to verify each SAS IT Resource Management adapter. This information is not an exclusive list of all data source versions that are supported. Earlier versions of some data source software might be supported and future releases of some data source software might be compatible and supported by adapters for SAS IT Resource Management. For information about supporting a data source version that is not

included in the following table, contact your SAS Installation Representative or on-site SAS support personnel.

Supported Adapters

The name of an adapter represents the data source from which SAS IT Resource Management gathers raw data.

Note: All SAS IT Resource Management adapters are supported on all SAS IT Resource Management operating environments unless there is a technical limitation that is associated with the raw data (log file) sources. For example, if the raw data resides on a device that is not compatible with the architecture of the host system on which the SAS IT Resource Management server is installed, then that raw data file might not be usable as input unless further processing is performed.

Table A2.1 Table of Supported Adapters

Adapter Name	Description	Latest Verified Version of the Data Source
ASG TMON2CIC	<p>ASG TMON for CICS TS for z/OS (formerly from Landmark) monitors the CICS Transaction Server (TS) and provides resource consumption measurements by tracing each transaction's performance by CICS event and related unit of work.</p> <p>Note: CICS data can be produced by either ASG TMONCICS or as a subtype of SMF.</p>	MXG 29.04*
ASG TMONDB2	<p>ASG TMON for DB2 (formerly from Landmark) works with the SQL Analyzer to manage DB2 applications and critical DB2 resources by providing a single view of DB2 data-sharing group performance. ASG TMON for DB2 tracks DB2 buffer pools, lock contention and serialized resource usage, transaction statistics from ASG TMON for CICS to ASG TMON for DB2, and dynamic and static SQL calls. ASG TMON for DB2 provides a view for all performance statistics by plan, package, and statement, including CPU time, I/O analysis, lock activity, and buffer pool utilization by SQL statement.</p>	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
BMC Mainview IMS	BMC Mainview for IMS (formerly IMF) supplies transaction information such as CPU usage, number of I/Os, and response time. When you collect BMC Mainview for IMS data instead of IBM IMS data, BMC recommends that you collect a duration of one hour of consistently high activity, starting at the beginning and ending at the end of an RMF interval.	MXG 29.04*
BMC Perf Mgr	BMC Performance Manager for Servers (formerly BMC PATROL) collects UNIX and Windows enterprise system management data.	v3.7.30
CA TMS	The Computer Associates CA 1 Tape Management product (formerly TMS) controls and protects tape data sets and volumes in z/OS environments.	MXG 29.04*
DT Perf Sentry	Demand Technology Performance Sentry (formerly NTSMF) collects data from Windows systems and for Windows server applications such as Exchange, SQL Server, and IIS. SAS IT Resource Management also supports a separate Demand Technology Performance Sentry adapter that is z/OS based. The Demand Technology Performance Sentry with MXG adapter that is z/OS based is for raw data sources that are supported through MXG.	v3.1.4
DT Perf Sentry with MXG	Demand Technology Performance Sentry (formerly NTSMF) with MXG collects data from Windows systems and for Windows server applications such as Exchange, SQL Server, and IIS. (SAS IT Resource Management also supports a separate Demand Tech Perf Sentry adapter that is not MXG based.)	MXG 29.04*
HP Perf Agent	HP Performance Agent (formerly OpenView Performance Agent) captures enterprise systems management measurement data for IT resources from Windows systems and many UNIX variants such as HP-UX, Sun OS/Solaris, IBM AIX, Tru 64 UNIX, and Linux.	v5.0

Adapter Name	Description	Latest Verified Version of the Data Source
HP Reporter	HP Reporter software (formerly OpenView Reporter) captures HP Performance Agent measurement data that is stored in a relational database via SAS/ACCESS to Oracle or ODBC.	v3.70
IBM AS400	<p>IBM OS/400 Collection Services</p> <p><i>Note:</i> The staging code that is required for the IBM AS400 data source is not automated by a SAS IT Resource Management adapter. Instead, SAS IT Resource Management supplies a User-written Staging transformation and templates that contain table and column metadata for this adapter for use with user-written staging code.</p>	MXG 29.04*
IBM DCOLLECT	<p>IBM DFSMS Data Collection Facility obtains the following Mainframe Storage System data:</p> <ul style="list-style-type: none"> • active data sets • VSAM data set information • volumes • inactive data • migrated data sets • backed-up data sets • capacity-planning data such as the following: <ul style="list-style-type: none"> - DASD capacity planning: IBM DCOLLECT provides information and statistics for volumes that are managed by DFSMSHsm (ML0 and ML1). - tape capacity planning: IBM DCOLLECT provides statistics for tapes that are managed by DFSMSHsm. • SMS configuration information 	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
IBM EREP	<p>IBM Environmental Record Editing and Printing Program processes the error records from the MVS, VM, and VSE operating systems to produce formatted reports. These IBM EREP reports can show the status of the entire installation, an I/O subsystem, or an individual device, depending on which report you request. IBM EREP reports can vary in format, depending on the following report types:</p> <ul style="list-style-type: none"> • System summary reports show error data in summary form. • Trend reports show error data by daily totals. • Event history reports show error data in a time sequence by occurrence. 	MXG 29.04*
IBM IMS	<p>IBM Information Management System (IMS) is the transaction and hierarchical database management system from IBM.</p> <p><i>Note:</i> The staging code that is required for the IBM IMS data source is not automated by a SAS IT Resource Management adapter. Instead, SAS IT Resource Management supplies a User-written Staging transformation and templates that contain table and column metadata for this adapter for use with user-written staging code.</p>	MXG 29.04*

Adapter Name	Description	Latest Verified Version of the Data Source
IBM SMF	<p>IBM System Management Facility collects and records system and job-related information that is used for the following purposes:</p> <ul style="list-style-type: none"> • billing users • reporting reliability • analyzing the configuration • scheduling jobs • summarizing direct access to volume activity • evaluating data set activity • profiling system resource use • maintaining system security <p>Note: RMF is a subtype of SMF records. IBM RMF is the IBM strategic product for z/OS performance measurement and management. IBM RMF is the base product that collects performance data for z/OS and Sysplex environments to monitor system performance. Because IBM RMF data is a subtype of IBM SMF data, the IBM SMF adapter can be used with IBM RMF data without requiring additional customizations.</p>	MXG 29.04*
IBM TPF	<p>The IBM Transaction Processing Facility operating system works with application programs to process transactions in a real-time environment. The IBM TPF system is designed for businesses and organizations that have large networks and high volumes of online transactions.</p>	MXG 29.04*
IBM VMMON	<p>IBM VM Monitor Collection product collects performance information that is associated with VM, including measurements for user activity, processor storage, I/O, and applications.</p>	MXG 29.04*
MS SCOM	<p>Microsoft System Center Operations Manager monitors Windows systems, Microsoft software, and applications to provide a view of the health of an organization's Windows environment.</p>	v2007

Adapter Name	Description	Latest Verified Version of the Data Source
SAP ERP	<p>Enterprise Resource Planning (ERP) Software from SAP (formerly SAP R/3), via the SAS IT Management Adapter for SAP, collects SAP ERP and SAP Business Warehouse (BW) workload performance measurements.</p> <ul style="list-style-type: none"> For SAP ERP, measurements are read from the SAP Statistic File (stat file). For SAP BW, the measurements that communicate the amount of computer resources that were consumed to populate and query BW cubes are read. These measurements are then used to supplement BW server performance statistics. 	v7.0
SAR	<p>System Activity Reporter is a logging mechanism that is native to most UNIX and Linux variants. SAR captures the contents of cumulative system activity counters.</p>	<p>Dependent on the operating system.</p> <p>For information about specific operation system release versions, see “Accessing Raw Data for SAR” in Appendix 2, “Data Sources Supported by SAS IT Resource Management Adapters” in SAS IT Resource Management: Administrator’s Guide.</p>
SNMP	<p>Simple Network Management Protocol is a component of the Internet Protocol Suite.</p> <p>The SNMP adapter also supports gathering SNMP data from round-robin databases (RRDs) that are used in many network management tools.</p>	Not applicable

Adapter Name	Description	Latest Verified Version of the Data Source
VMware vCenter	VMware vCenter Server is an operational console for VMware and includes the VMware vCenter database that serves as the raw data source for this SAS IT Resource Management adapter. The vCenter database contains VMware performance and utilization measurements for both the virtual machines and the physical servers on which they run, including detailed CPU, memory, storage, and I/O metrics. The vCenter relational database is accessed via SAS/ACCESS to Oracle, DB2, or ODBC.	v4.0
VMware Data Acquisition	<p>VMware Data Acquisition is a transformation that can extract topology and configuration information from the VMware vCenter database more frequently than once a day. The vCenter relational database is accessed via SAS/ACCESS to Oracle, DB2, or ODBC.</p> <p><i>Note:</i> This is not an adapter. However, a staging transformation is available for this source.</p>	v4.0
Web Log	The Web Server Log adapter reads and analyzes Web logs that conform to the Common and W3C Extended Log Formats (CLF and ELF.)	Not applicable

* SAS IT Resource Management uses MXG for this adapter.

Note: The specific software versions of the raw data sources in this table have been used to verify each SAS IT Resource Management 3.21 adapter. This information is not an exclusive list of all data source versions that are supported. Earlier versions of some data source software might be supported. Future releases of some data source software might be compatible and supported by adapters for SAS IT Resource Management 3.21. For information about supporting a data source version that is not included in the previous table, contact your SAS Installation Representative or on-site SAS support personnel.

For data sources (and for raw data sources) that are user-written, staging code is supported by the wizards of SAS Data Integration Studio, and the transformations that are supplied by SAS IT Resource Management. The New Table wizard of SAS Data Integration Studio can create the transformations that can stage raw data from any input source.

Accessing Raw Data for MXG

Preparing MXG Data for SAS IT Resource Management

MXG supports the following data sources:

- ASG TMON2CIC
- ASG TMONDB2
- BMC Mainview IMS
- CA TMS
- DT Perf Sentry with MXG
- IBM DCOLLECT
- IBM EREP
- IBM SMF
- IBM TPF
- IBM VMMON

For information about the IBM SMF performance areas that are collected by MXG, see [“SMF Adapter - RMF Domain Category” on page 479](#) and [“SMF Adapter - Jobs, TSO, and OMVS Domain Categories” on page 483](#).

Extracting Data from MXG

MXG creates SAS data sets for each of the data sources that it supports. These data sets can be accessed directly by the adapter staging code.

Accessing Raw Data for HP Perf Agent

Preparing HP Perf Agent Data for Staging

HP Performance Agent software, formerly OpenView Performance Agent, captures enterprise systems management measurement data from Windows systems and many UNIX variants such as HP-UX, Sun OS/Solaris, IBM AIX, Tru 64 UNIX, and Linux.

For information about the performance areas that are collected by HP Perf Agent, see [“HP Perf Agent and HP Reporter Adapters” on page 474](#).

Extracting Data from HP Perf Agent

The log files that contain the performance data can be exported by using **extract**, an HP Performance Agent program. The **extract** program writes the performance data in binary format. This binary file is then used as input to SAS IT Resource Management.

The following list is a partial list of the options and arguments that can be used to specify how to run the **extract** command. For the complete list of options and arguments that are available, see the *HP OpenView Performance Agent User's Manual*.

- xp
exports the extracted data to the specified output file.
- d
extracts all the data from a single day.
- f *output-filename*
specifies that the output is to be written to an output file that can be read directly into SAS IT Resource Management.
- gapdcnt
specifies the types of data to extract.
- i
specifies the extraction of logical system data for the host and any associated guest machines present in the virtual environment. If -i is not specified, the extraction can still work, but no logical system data is extracted.

Note: The -i option applies to the designated host machine in a virtual operating environment. At this time, the -i option applies to the following three platforms: HP-UX (virtual machine configuration), AIX (LPARs or logical partition configuration), and Linux (VMWARE ESX configuration).
- r *export-template-filename*
specifies the name of the export template file. This file governs the format that is used for the exported data.

Note: Only data that was logged by HP Perf Agent for Windows NT/2000 can be configured and modified for use by export templates. Export templates cannot be used on any other platforms.
- v
generates verbose output report formats.

Verbose mode enables the echoing of commands and command results.

Example 1

The following command extracts all the data for yesterday (d-1). It exports the data to the output file that is named data.bin, using the template in a file that is named reptfile.mwr. The purge parameter specifies that any existing file is purged in order to make room for the new file. The gapdcnt option specifies that the following types of data are to be extracted and exported: global detail, application detail, process detail, disk device detail, configuration detail, netif detail, and transaction detail.

```
extract -xp d-1 -gapdcnt -r reptfile.mwr -f data.bin,purge -v
```

The log file that is created by HP Performance Agent software can be input directly into SAS IT Resource Management. Alternatively, you can store multiple log files in a directory that can be input to SAS IT Resource Management.

It is advisable to put the **extract** command into a shell, bat, or Perl script that runs once a night by means of a scheduling tool. (Perl is an open-source programming language that can be used for data manipulation tasks.) Cron can be used for scheduling on UNIX systems, and Task Manager can be used on Windows operating systems. The extracted files from each server can then be transferred to a centralized location to be

read by SAS IT Resource Management using the file transfer utility of your choice. Delete the older raw data files to preserve disk space.

Example 2

The following shell script runs on Linux and extracts yesterday's data into files that have date extensions as part of the name.

```
$ cat /var/opt/perf/datafiles/itrm-extract/extract.sh
# shell script to do HPOV extracts
/opt/perf/bin/extract -xp d-1 -gapdcnt \
    -l /var/opt/perf/datafiles/logglob
\    -r /var/opt/perf/reptfile.mwr \
    -f /var/opt/perf/datafiles/itrm-extract/server1-'date +%y%m%d'.bin, purge -v
```

The output file that is generated by the preceding shell script is of the form **server1-yyymmdd.bin**. (If yesterday's date is November 26, 2007, then the name of the output file is **server1-071126.bin**). The output file is placed in the directory that is located at: **/var/opt/perf/datafiles/itrm-extract**.

For more information about generating HP Perf Agent data files, see the *HP OpenView Performance Agent User's Manual*.

Accessing Raw Data for DT Perf Sentry

Preparing DT Perf Sentry Data for SAS IT Resource Management

SAS IT Resource Management provides both a native DT Perf Sentry adapter and an adapter based on MXG (DT Perf Sentry with MXG). IT Resource Management supplies aggregations with the native DT Perf Sentry adapter. Therefore, using the native DT Perf Sentry adapter is recommended.

Note: Demand Technology Performance Sentry software is included in your SAS Software Depot for IT Resource Management when the server tier is licensed on a Windows system. The executable software that installs Demand Technology Performance Sentry is in the **products\perf Sentry_92110_prt_xx_sp0_1** folder. The Performance Sentry license installation enables you to manage up to 100 Windows systems and is set to expire 180 days after installation. At that time, you must contact Demand Technology to arrange to continue the use of this software at your site. For information about this software, see <http://www.demandtech.com>.

The Demand Technology Performance Sentry product (formerly NTSMF) collects a variety of Windows NT performance objects and their associated counters. DT Perf Sentry gathers and calculates object counter values that are based on the counter type. It then writes them on a periodic basis to an ASCII-formatted, comma-delimited, or tab-delimited data file. This file is input directly into SAS IT Resource Management.

The following requirements apply to the collection of DT Perf Sentry data:

- For Windows and UNIX operating environments, if you transfer files across operating systems, make sure that the files retain their ASCII format. (If you want to process your data on z/OS, then the data must be converted from ASCII format to EBCDIC. For information, see “[Preparing DT Perf Sentry Data for Processing on z/OS](#)”.)

- Discovery records must be collected in DT Perf Sentry smf log files and must be written for all types of data records.
- If your version of Performance Sentry Administration enables you to select which Discovery record types are collected, select the type 5 records. However, record types 6 and 7 can also be selected because these records do not add any significant overhead in terms of space or performance.
- Compressed and uncompressed records can be processed.
- The delimiter that is used in DT Perf Sentry data must be a comma or a tab character.
- _Total records must not be written to the data file.
- DT Perf Sentry filters must not be used.

For information about the performance areas that are collected by DT Perf Sentry, see “DT Perf Sentry Adapter - DT Perf Sentry Server-Based Domain Categories” on page 470 and “DT Perf Sentry Adapter - DT Perf Sentry Exchange Domain Category” on page 471.

Preparing DT Perf Sentry Data for Processing on z/OS

In addition to the preceding requirements, if you want to process DT Perf Sentry logs on z/OS, you must allocate a PDS with the following DCB attributes on z/OS:

DSORG=PO, RECFM=VB, LRECL=256

Note: Accept the default for BLKSIZE.

Extracting Data from DT Perf Sentry

The log file that is created by the DT Perf Sentry agent program can be input directly into SAS IT Resource Management 3.1.1 and later. It contains two types of records:

- discovery records that contain the format of the data records
- data Interval records that contain the counters for the performance areas that are being measured

Each DT Perf Sentry smf file should be maintained as a separate smf file. Concatenating smf files is not recommended. The DT Perf Sentry adapter accepts the pathname to either a raw data file or a directory that contains the input files for the adapter.

If a directory is specified, the directory cannot contain subdirectories or any files other than valid DT Perf Sentry raw data files. DT Perf Sentry does not manage the transfer of the data to a central location.

- On the UNIX and Windows operating environments, place the multiple smf files into a single directory where they can be input to the staging transformation in SAS IT Resource Management. No other files or subdirectories should be in this directory. Use the file management utility of your choice to accomplish this transfer of data.
- On the z/OS operating environment, use the File Transfer Protocol (FTP) to transmit each smf log into a unique PDS member, converting it from ASCII to EBCDIC as it is transmitted. By specifying the PDS name in the input location, each member is processed.

For more information, see the *Demand Technology Performance SeNtry User's Guide*.

Notes about DT Perf Sentry

After installation, DT Perf Sentry logs the performance data based on default parameter values that are stored in the Data Collection Set (DCS) Definition. The parameter values determine what data elements to collect and how frequently to collect them. For more information about the Data Collection Set Definition, see the "Performance SeNTRY Data Administration" chapter of the *Demand Technology Performance SeNTRY User's Guide*.

For best results, use the following two collection sets within DT Perf Sentry: ITRM Windows Server and ITRM Exchange. These collection sets are designed to correspond to the aggregations that SAS IT Resource Management provides for the DT Perf Sentry adapter. However, you do not have to use the ITRM Windows Server and ITRM Exchange collection sets in order to load DT Perf Sentry data into SAS IT Resource Management IT data marts. You can modify the objects and counters within the DCSs as needed.

Note: You can also access the ITRM Windows Server and ITRM Exchange collection sets from the SAS IT Resource Management Web site at: <http://support.sas.com/documentation/onlinedoc/DTPerfSentry/index.html>.

These DCS files are included wherever the IT Resource Management server tier is installed. You should move these files to the system where you installed DT Perf Sentry. Then, during the setup of the Demand Technology Performance Sentry product on the Windows system, instruct the product to use these two files. This action enables the correct set of metrics to be recorded on the DT Perf Sentry log files.

- By default, if the IT Resource Management server is installed on z/OS, a PDS called ITRM.CPMISC is automatically created with the ITRMEXCH and ITRMSRVR members. Use the File Transfer Protocol (FTP) to transmit the ITRMEXCH and ITRMSRVR members of the ITRM.CPMISC installation data set to the Windows system where Demand Technology Performance Sentry software is installed. (These members of the PDS are plain text fields and do not have to be transmitted in binary mode.)
- By default, if the IT Resource Management server is installed on UNIX, the DCS files are called ITRMExch.DCS and ITRMSrvr.DCS. The files can be located at `/SASHome/SASFoundation/9.3/misc/itmsmvadata`. Use the File Transfer Protocol (FTP) to transmit the two DCS files to the Windows system where Demand Technology Performance Sentry software is installed.
- By default, if the IT Resource Management server is installed on Windows, the DCS files are called ITRMExch.DCS and ITRMSrvr.DCS. The files can be located at `C:\Program Files\SASHome\SASFoundation\9.3\itmsmvadata\sasmisc`. You can copy and paste the files from this location to the Windows system where Demand Technology Performance Sentry software is installed.

Accessing the ITRM Windows Server DCS

To access the IT Resource Management Windows Server DCS, perform the following steps:

1. Download the ITRM Windows Server data collection set from <http://support.sas.com/documentation/onlinedoc/itsv/32/dtps.html>, and place it in a location that is accessible to the Demand Technology Performance Sentry DCS Administration client.

2. Open the Sentry Administration client.
3. In the DCS Administration window, select **File** ⇒ **Import DCS**.
4. Provide the location where the ITRM Windows Server data collection set was downloaded. The ITRM Windows Server data collection set should appear in the list of Available Data Collection Sets.

Note: The DCS files are also available wherever the IT Resource Management server tier is installed. If you want to use these files, you should move them to the system where your DT Perf Sentry is installed. For information, see the preceding “[Notes about DT Perf Sentry](#)”.

Accessing the ITRM Exchange DCS

To access the ITRM Exchange DCS, perform the following steps:

1. Download the ITRM Exchange data collection set from <http://support.sas.com/documentation/onlinedoc/itsv/32/dtps.html>, and place it in a location that is accessible to the Demand Technology Performance Sentry DCS Administration client.
2. Open the Sentry Administration client.
3. In the DCS Administration window, select **File** ⇒ **Import DCS**.
4. Provide the location where the ITRM Exchange data collection set was downloaded. The ITRM Exchange data collection set should appear in the list of Available Data Collection Sets.

Note: The DCS files are also available wherever the IT Resource Management server tier is installed. If you want to use these files, you should move them to the system where your DT Perf Sentry is installed. For information, see the preceding “[Notes about DT Perf Sentry](#)”.

Accessing Raw Data for BMC Perf Mgr

Preparing BMC Perf Mgr Data for SAS IT Resource Management

BMC Performance Manager for Servers data is collected by BMC Performance Manager, which was formerly BMC Patrol. BMC Perf Mgr data can be collected from the following operating systems:

- Windows NT
- UNIX
- SAP ERP

The BMC Perf Mgr Agent maintains parameter history data about the systems that it monitors. For more information, see your BMC Perf Mgr documentation.

BMC Perf Mgr enables each metric to be sampled at its own interval. This interval can be set by the BMC Perf Mgr administrator. SAS IT Resource Management requires that the sample rates be specified on minute boundaries. (However, you can also specify a sample rate of 30 seconds.)

SAS IT Resource Management provides support for three Knowledge Modules:

- Windows NT Knowledge Module
- UNIX Knowledge Module
- SAP ERP Knowledge Module Version 2.2

For information about the performance areas that are collected by BMC Perf Mgr, see [“BMC Perf Mgr Adapter - BMC Perf Mgr Domain Category” on page 472](#).

Extracting Data from BMC Perf Mgr

A raw data file of BMC Perf Mgr history data can be generated in two ways:

- extracting data from the BMC Perf Mgr History Knowledge Module

The history knowledge module organizes the collection of the history data and ensures that it is sent to a central server from which it can be extracted. For more information about this method, see the documentation from BMC Perf Mgr about the History Loader Knowledge Module.

- extracting data by using the **dump_hist.exe** command

This command extracts the same history data that is extracted from the History Knowledge Module. However, it does not manage the transferal of the data to a central location. (This option is useful if you prefer writing your own scripts to control the extraction and transferal of the data to a central location.)

Note: When transferring the history data across platforms, ensure that the data retains its ASCII text format.

Although these two methods produce slightly different output, both can be processed by the staging transformation of SAS IT Resource Management. The BMC Perf Mgr Operator Console retrieves the historical data that is stored by the Agent. The **dump_hist** line command copies the parameter history data that is maintained by the BMC Perf Mgr Agents. The BMC Perf Mgr Agent Reference Manual contains more detailed information about the **dump_hist** command. (The **dump_hist.exe** program is provided as part of the BMC Performance Manager software.)

The following command copies parameter history data for one day for the file that is indicated by *filename*. The command uses start (-s) and end (-e) parameters to specify the date and time of data that is to be copied. The format of these parameters is **MMddhhmm [yyyy]**, where

- MM= month
- dd= day
- hh= hour
- mm= minute
- yyyy= year, which is optional

The following code copies the full day of data from July 18, 2007.

```
- dump_hist -s 071800002007 -e 071823592007 > filename
```

The following example shows the format of the text file that is created by the preceding **dump_hist** command. This file can be input to the staging transformation of SAS IT Resource Management.

```
machinename/NT_CPU.CPU_0/CPUprcrUserTimePercent
```

```

Thu Jul 18 10:00:57 2007 26.981
Thu Jul 18 10:01:58 2007 5.35963
Thu Jul 18 10:02:58 2007 0.598205
Thu Jul 18 10:03:58 2007 0.333915
machinename/NT_CPU.CPU_0/CPUprcrPrivTimePercent
Thu Jul 18 10:00:57 2007 61.0279
Thu Jul 18 10:01:58 2007 1.20528
Thu Jul 18 10:02:58 2007 1.56053
Thu Jul 18 10:03:58 2007 1.05312
machinename/NT_SYSTEM.NT_SYSTEM/SYSsysTotalProcTimePercent
Thu Jul 18 10:00:57 2007 88.013
Thu Jul 18 10:01:58 2007 6.56211
Thu Jul 18 10:02:58 2007 2.1812
Thu Jul 18 10:03:58 2007 1.36592

```

To dump data for a single server, add the `-host<hostname>` parameter to the previous command. The command then dumps a full day of data for the single server that is identified by `<hostname>`.

For additional information about the syntax of the `dump_hist`, see the *PATROL® Script Language Reference Manual*.

Notes about BMC Perf Mgr

When staging BMC Perf Mgr data on z/OS, allocate the z/OS file with an LRECL= 200. (BMC Perf Mgr data typically has variable length records, but they do not exceed 200 bytes in length.)

Processing BMC Perf Mgr history data into a staged table might require the following two issues to be resolved:

- different sample rates are collected for each metric

Two metrics 'A' and 'B' are not necessarily sampled at the same rate. Metric 'A' might be sampled at 1-minute intervals, and 'B' might be sampled at 5-minute intervals. To combine these two metrics into the same observation in a table is invalid because each value should eventually be weighted by the duration variable. To resolve this problem, the staging code of SAS IT Resource Management includes a variable in each BMC Perf Mgr table called DURGRP. DURGRP is a string that represents the duration group to which a metric belongs. In this example, metric 'A,' which is sampled every minute, is included in the observation with a DURGRP value of 60 (60 seconds). Metric 'B' is included in an observation with a DURGRP of 300 (300 seconds).

The DURGRP variable is used only during staging to ensure that the metrics are reduced and summarized by their respective duration values (assuming that they are weighted by DURATION).

If numerous null values are present in each observation, then the BMC Perf Mgr data in the SAS IT Resource Management tables might look unusual. The number of DURGRPs and null values depends on the number of different sample rates that are applied to metrics that belong to the same table.

- Datetime Stamps of Samples Are Not Exactly Aligned

In this example, two metrics 'A' and 'B' are both sampled at 1-minute intervals. In the following example history data, the first sample occurred at x for both metrics. However, in the second sample, the datetime stamps differ by a second, with 'B' being sampled later than 'A'. The first sample for each metric is combined into a

single observation because the duration and datetime stamps are the same. However, this is not the case for the second sample.

```
machinename/NT_CPU.CPU_0/A
Thu Jul 18 10:00:57 2007 26.981
Thu Jul 18 10:01:58 2007 5.35963
machinename/NT_CPU.CPU_0/B
Thu Jul 18 10:00:57 2007 61.0279
Thu Jul 18 10:01:57 2007 1.20528
```

During the staging of the raw data, SAS IT Resource Management detects that this second sample has related datetime values and collapses the data into one observation. The result is that the data in the staged table is less sparse. However, the datetime and duration values are only near approximations.

Accessing Raw Data for SAR

Preparing SAR Data for SAS IT Resource Management

System Activity Reporter (SAR) data can be collected from the following operating systems:

- HP-UX 10.2 and later
- IBM AIX 3.2 and later
- SunOS Solaris 5.8 and later
- Linux (Use Sysstat 4.0.1 or Sysstat 5.0.0 and later.)

SAR performance data can be set up to write continuously to a binary file by means of the **sapd** command. At appropriate intervals, the **sar** command can be used to write the binary file to an ASCII file. The ASCII file is input to a staging transformation of SAS IT Resource Management.

For information about the performance areas that are collected by SAR, see [“SAR Adapter - SAR Domain Category” on page 478](#).

Extracting Data from SAR

The following **sapd** and **sar** commands and shell scripts that collect SAR data are usually available in the `/usr/lib/sa` directory:

- The **sapd** command collects performance data into a binary file.

The `sa1` shell script is a wrapper for the **sapd** command and can be scheduled to run automatically through cron.

By default, the binary file that is created by the `sa1` shell script is called `/var/adm/sa/sa$DATE` (where `$DATE` is the current date). To save the binary data in a different location, modify the `sa1` shell script accordingly.

The following cron entry shows an example of running the `sa1` shell script to collect performance data every five minutes. The `sa1` shell script is run every 20 minutes by cron, and within the `sa1` shell script, the **sapd** command collects data for four 300 second (that is, 5-minute) intervals:

```
0,20,40 * * * * /usr/lib/sa/sa1 300 4
```

- The **sar** command extracts the data from the binary file that was created by the **sapd** command and creates an ASCII file that can be read by the SAR staging transformation of SAS IT Resource Management 3.1.1.

The **sa2** shell script is a wrapper for the **sar** command, and can be scheduled to run automatically through cron.

By default, the binary file is **/var/adm/sa/sa\$DATE**. To change the location of the binary file or the location of the ASCII file, modify the **sa2** shell script accordingly.

The following cron entry shows an example of running the **sa2** shell script on Linux to create ASCII files once a day at 12:55AM:

```
55 0 * * * /usr/lib/sa/sa2 -A
```

Note: The switches for the **sar** command can be passed as parameters to the **sa2** shell script.

Using the sar Command

The **sar** command can be used to collect the following types of data:

- global data (data for the entire system)
- processor level data (data for each processor on the system)

If processor level data is collected, it can be staged in the SARCPUB table. Each UNIX environment has its own set of options that can be specified with the **sar** command. SAS IT Resource Management recommends setting the following switches:

- for HP-UX operating environments:

To collect global data as well as per-processor data on HP-UX, you must create two files by executing two **sar** commands.

- To collect global data, execute the following command:

```
sar -A -f /var/adm/sa/sa$DATE > /var/adm/sa/sar$DATE_global
```

- To collect processor level data, execute the following command:

```
sar -Mu -f /var/adm/sa/sa$DATE > /var/adm/sa/sar$DATE_processor
```

- for AIX operating environments:

To collect global data as well as per-processor data on AIX, you must create two files by executing two **sar** commands.

- To collect global data, execute the following command:

```
sar -bdkqrvy -f /var/adm/sa/sa$DATE > /var/adm/sa/sar$DATE_global
```

- To collect processor level data, execute the following command:

```
sar -A -P ALL -f /var/adm/sa/sa$DATE > /var/adm/sa/sar$DATE_processor
```

- for Linux operating environments:

To collect both global and processor level data, use the following **sar** command:

```
sar -A -f /var/adm/sa/sa$DATE > /var/adm/sa/sar$DATE
```

- for SunOS Solaris operating environments:

You cannot capture processor level data on SunOS Solaris. To collect global data, execute the following **sar** command:

```
sar -A -f /var/adm/sa/sa$DATE > /var/adm/sa/sar$DATE
```

All of the files that are created by the **sar** command should be placed in a single directory that is read by the staging transformation in SAS IT Resource Management 3.1.1. In the staging transformation, all the files in this directory are read and used to update the SAR data mart.

Note: For more information, see your system's documentation about the **sar** and **sapd** commands and the **sa1** and **sa2** shell scripts.

Accessing Raw Data for SNMP

Sources of SNMP Data

Simple Network Management Protocol (SNMP) data can be collected from the following sources:

- HP Network Node Manager (HP NNM)
- Round-Robin Database

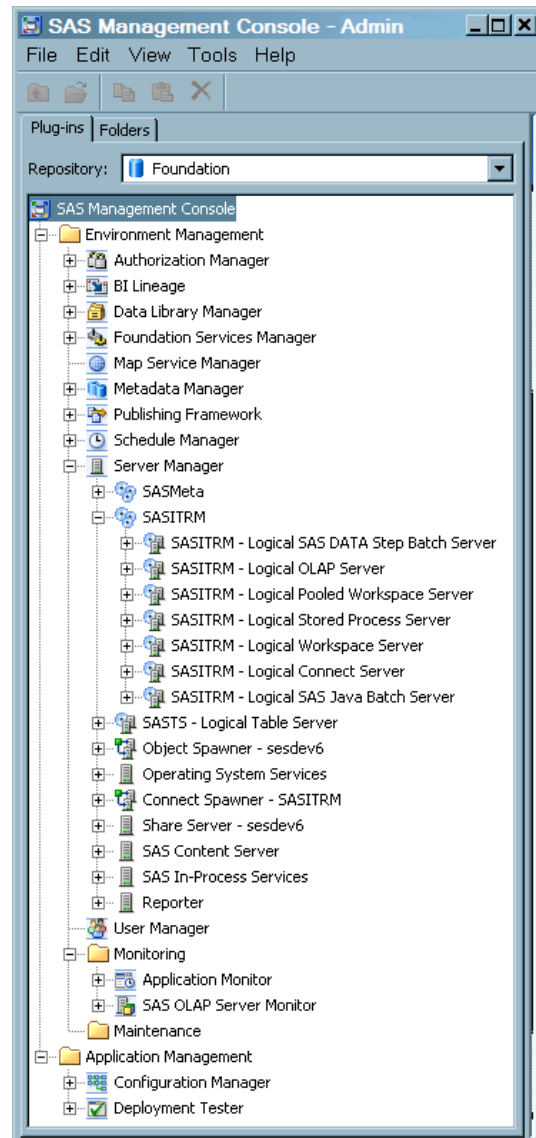
In order to read the raw data from either HP NNM or rrdtool, it is necessary to run Perl scripts as part of the staging process. This means that Perl must be installed on the machine that the SAS program runs on, that is, the SAS Workspace Server, the SAS batch server, or interactive SAS.

The Perl scripts that are used by the staging code are generated as part of the code generation of the staging job. This code can be viewed and modified as necessary in the **Code** screen of the job.

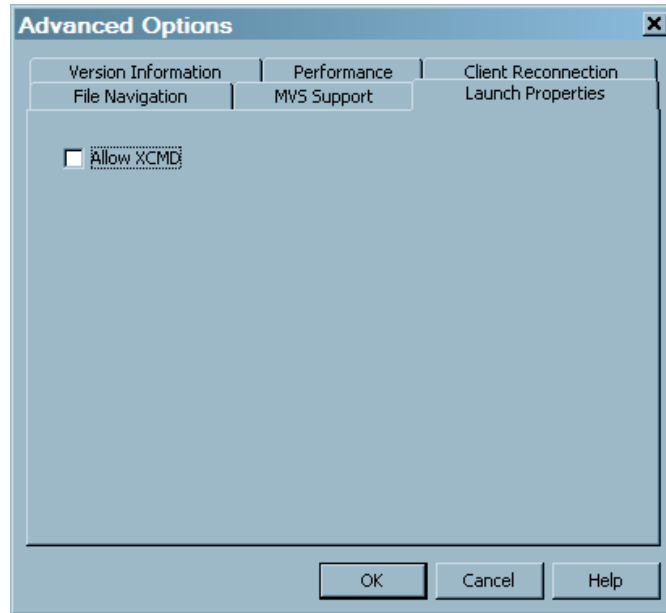
To run these scripts in batch mode, it is necessary to set the **xcmd** option for the Batch server. To do so, change **set USERMODS_OPTIONS=** to **set USERMODS_OPTIONS=XCMD**.

To run these scripts from within the SAS Workspace server, it is necessary to have the SAS option **XCMD** turned on. By default this option is off in the Workspace server. To turn the option on, perform the following steps:

1. Launch SAS Management Console. Log on as an administrator.
2. Expand the **Server Manager** in the left panel.
3. Expand **SASITRM**. Then expand **SASITRM-Logical Workspace Server**.

Display A2.1 Expanded Left Panel of SAS Management Console

4. Right-click **SASITRM-Workspace Server** and select **Properties**. The Workspace Server Properties dialog box appears.
5. Select the **Options** tab and click **Advanced Options**.
6. Select the **Launch Properties** tab to open the following dialog box.

Display A2.2 Launch Properties Tab of the Advanced Options Dialog Box

7. Select the **Allow XCMD** check box.
8. Click **OK** to close all the open dialog boxes.
9. Stop and then restart your Object Spawner service.

Windows Specifics

To stop your object spawner service, from the Start menu, select **Programs** ⇒ **SAS** ⇒ **SAS Configuration** ⇒ *<configuration-name>* ⇒ **Object Spawner** ⇒ **Stop**. To restart your object spawner service, from the Start menu, select **Programs** ⇒ **SAS** ⇒ **SAS Configuration** ⇒ *<configuration-name>* ⇒ **Object Spawner** ⇒ **Stop**.

UNIX Specifics

To stop and then restart your object spawner service, from the command line, change directories to SAS-config-dir/Levl/SASMain/ObjectSpawner. Stop the object spawner by issuing this command: `$./ObjectSpawner.sh stop`. When you receive a confirmation that the object spawner has stopped, start it again by issuing this command: `$./ObjectSpawner.sh start`. You should receive a confirmation that the object spawner has started.

z/OS Specifics

To stop and restart your object spawner service, stop and restart the z/OS started task associated with the object spawner. Follow the stop and restart procedures that are established at your site.

Extracting Data by Using HP NNM Tools

Versions of HP NNM that are earlier than HP NNMi 8.x provide two tools that are no longer provided with the HP NNMi 8.x and later versions of the product. However, if you have access to earlier releases of the HP NNM product, you can still use the commands that are discussed in the following section.

Versions of HP NNM that are earlier than i8.x provide two commands that enable you to extract SNMP data. The two commands are:

- `snmpCollect`, which gathers SNMP data into binary files
- `snmpColDump`, which reads the binary files and writes the data into ASCII files that can be processed by the SNMP staging transformation

The `snmpCollect` executable uses a file called `snmpCol.conf` to determine the metrics to collect and the devices from which to gather data.

You can create the `snmpCol.conf` file by using either of the following methods.

- Use the sample UNIX script (`snmpmkcf`) that SAS IT Resource Management provides. The sample script is located in the `sasmisc` folder. It reads the `snmphost` and the `snmpvars` files, which are also located in the `sasmisc` folder, and creates the `snmpCol.conf` file.
 - The `snmphost` file contains a list of the device names for which SNMP data is being gathered.
 - The `snmpvars` file contains a list of the SNMP metrics that are to be gathered from each device. By default, the values in this list are the metrics that are in the recommended staging tables.

Note: This method generally creates the `snmpCol.conf` file more quickly than using the NNM graphical user interface.

- Use the HP NNM graphical user interface to create `snmpCol.conf`. The user interface prompts you to select the metrics that you want to gather and the devices about which you want to gather data. Using your responses, it then creates the `snmpCol.conf` file.

When the `snmpCol.conf` file is created, you can run the `snmpCollect` command to gather data. You can run the collect as often as you choose. The command collects a set of files in the location that you specify. For each metric and interface, two files are generated:

- `metricName.interface-number`, that contains the raw data
- `metricName.interface-number!`, that contains the information about the metrics that are collected

For example, for the metric called `ifOctets`, and for interface 3, the `ifOctets.3` and `ifOctets.3!` files are generated. The SNMP staging code runs the `snmpColDump` command against these files to get the rawdata that is to be read and loaded into staged tables. In order to read the data, the `snmpColDump` executable requires both the file that contains the binary data and file that contains the information about the metrics that are collected.

TIP The binary data files continue to accumulate until they are deleted. For best results, these data files (whose names do not end with "!") should be moved to another location before staging. The `snmpCollect` command then creates a new set of binary files into which the data is to be collected. The files that contain the information about the metrics (whose names end with "!") should be copied (not moved) to the other location as well. This way, the `snmpCollect` command can continue to use these metrics as it collects the raw data. The staging code can then read the files from the location to which they were moved or copied.

Extracting Data by Using the Round-Robin Database Tool (RRDTOOL)

The SNMP adapter can read data from a round-robin database that was created with `rrdtool`.

Note: The round-robin database tool, `rrdtool`, is a freeware package that is available for download from Tobias Oetiker. This software is widely used, often in conjunction with Tobias Oetiker's MRTG package, which is a real-time monitor of data.

The SNMP adapter reads only those round-robin databases that contain one round-robin database for each combination of the device, the network interface to which the device is connected, and the metric. The data can be gathered at any interval, and any metrics and devices can be gathered.

In order to produce the correct set of round-robin databases, SAS IT Resource Management provides a sample Perl script for data collection and storage in the round-robin database. This script, called `snmpuprd.pl`, is available in the `sasmisc` folder. The section at the top of the script enables you to specify modifications that control how the data is collected. The Perl script uses the `snmphost` and `snmpvars` files, which are located in the `sasmisc` folder, to determine which metrics to gather from your devices. It also uses the `snmpwalk` command to collect the snmp data. When you have the script set up correctly, you can use your operating system's scheduling facility to run the script at the appropriate interval.

Note: You can backload data from the `rrdtool`. To so, specify the **Number of days to load** option in the **Staging Parameters** property tab of the SNMP staging transformation. (This option is available only if the chosen access command is `RRDTool`.)

Notes about the SNMPWALK Command

The `snmpwalk` command is an executable that is available as part of the Net-SNMP package. This package is available as freeware that you can download from the Sourceforge.net Web site. The `snmpwalk` command provides an easy way to read SNMP data from specified devices.

The SNMP adapter uses `snmpwalk` in two places:

- **Using `rrdtool`:** If you choose to store your rawdata using `rrdtool`, the script that is provided with SAS IT Resource Management uses `snmpwalk` to gather the data to put in the round-robin database.
- **Staging the data:** The SNMP staging transformation provides the ability to collect the non-numeric data at staging time instead of at the normal regular intervals that the rest of the SNMP data is collected. This can be useful for character data that does not change frequently. If you want to collect character data and are using `rrdtool` to gather your rawdata, then you must collect it at staging time, because `rrdtool` does not store character data. If you are using HP NNM, you can set it up to gather the character data at regular intervals along with the rest of the data, or you can choose to collect it at staging time.

If you want to use the `snmpwalk` method, go to the **Staging Parameters** tab of the SNMP transformation's property window. On the **Use `snmpwalk` to gather character data** option, select **Yes**. Then fill in values for the following options:

- **`snmpwalk` Executable:** This is the path and exe file that contains the `snmpwalk` command.
- **HostFile for `snmpwalk`:** This is the path and file that contains a list of the devices to poll for character data.
- **Community value for `snmpwalk`:** This option is usually "public." Check with your SNMP administrator to ascertain the correct value for your site.

If you choose to use `snmpwalk`, then a script is generated and executes when the staging code runs. This script executes the `snmpwalk` command in order to gather the SNMP data for specified character metrics and adds that data directly to the staged tables.

Working with the `ITMS_SNMP_DeviceType` Lookup Table

In order to provide better reporting, the SNMP adapter provides a SAS data table that enables you to indicate the type of device for each of the devices for which you are gathering data. This table is located by default in the SASHELP library, and contains no data. To use this table to define your device types, perform the following steps from the SAS IT Resource Management client:

1. Locate, or create, the IT data mart that you want to use for staging SNMP data.
2. Locate the **Administrative** folder. Right-click the library within that folder. (The library is called **Adminnnnn**, where *nnnn* is a random number that was generated when the IT data mart was created.)

Select **View Libname**.

3. Select the entire LIBNAME statement in the displayed box. Then right-click the LIBNAME statement and select **Copy**.
4. Launch an interactive session of SAS. Use the **Paste** function to paste the LIBNAME statement into an Editor window.
5. Submit the LIBNAME statement.
6. In the Explorer window, locate the SAS data table called `Itms_snmp_devicetype` by drilling down into the SASHELP library.
7. Use the **Copy** and **Paste** functions to copy this table to the **Adminnnnn** library.
8. Now you can edit the table that you copied to **Adminnnnn** library. To do so, use `VIEWTABLE` to add and update rows. (To access `VIEWTABLE`, double-click the table in the Explorer window.)

The value in the **Device** column should be the name of the device, and **deviceType** is the type of device.

- The name of the device should match the device for which SNMP is gathering data. Make sure that the device names use the same case.
- **DeviceType** can be any character string that helps identify the device. For example, "Router," "Switch," or "Server" would be reasonable choices for device Type.

You can also update this table programmatically as long as the table has the **device** and **deviceType** columns.

Once this table is populated, when you run the staging code, the **deviceType** column in all the SNMP staging tables will have the correct value from the lookup table. If a device is not in the deviceType table, then **deviceType** is set to "Unknown."

If your devices are named using a convention (for example, the device type is always the first part of the name, like `router_2123`), you can create the **deviceType** value programmatically. To do so, change the expression for the **deviceType** column in the staged table metadata. By default, the expression for the **deviceType** column is:

```
put(device, ?? $deviceType.)
```

You can change this to a different expression. An example is:

```
upcase(scan(device, 1, "_"))
```

Or, you can use standard SAS DATA step code to create an expression. An example is:

```
if (device eq: "R") then
    deviceType = "Router";
else if (device eq: "W") then
    deviceType = "Switch";
else if (device eq: "S") then
    deviceType = "Server";
else
    deviceType = "Unknown";
```

Working with Management Information Base (MIB) Definitions for SNMP Staged Tables

The SAS macro %rmmssnmp enables the SNMP adapter to read Management Information Base (MIB) definitions and create the metadata for SAS IT Resource Management staged tables. This macro enables you to generate staged data from any MIBs that you have. The SAS macro is available in the SAS autocall macro library so that you can use it from any SAS session.

To use the macro, you must first create a job with an SNMP transformation in SAS IT Resource Management. To do so, perform the following steps:

1. In the SAS IT Resource Management client, navigate to the IT Data Marts tree.
2. Navigate to an existing IT data mart, or create a new one. (To create a new IT data mart, from the menu bar, select **New** ⇒ **IT Resource Management** ⇒ **IT Data Mart**. The IT Data Mart wizard opens and prompts you to enter information to specify your new IT data mart.)
3. Create a folder in the IT data mart for your new job and tables.
4. In the new folder, create a new job. (To create a new job, select **New** ⇒ **Job**. Then enter the information to specify your new job.)
5. In the process flow diagram, drag the SNMP transformation from the **ITRM Transformations** folder of the Transformations tree and drop it onto the job.
6. Save and close the job.
7. Select the job in the tree and in the Basic Properties window, scroll to the **Metadata ID** property, and make a note of the value. This is the job ID that you need to run the macro in the next set of steps.

After the job is created, you can run the %rmmssnmp macro. To do so, perform the following steps:

1. Invoke an interactive SAS session.
2. Specify the parameters for the %rmmssnmp macro. For information, see [“Introduction to the Macros in SAS IT Resource Management” on page 510](#).
3. Submit the macro call to read the MIB definitions. The macro can generate a report of the staged tables that would be created from this definition, or it can create the staged table metadata.

Accessing Raw Data for Web Log

Preparing Web Log Data for SAS IT Resource Management

SAS IT Resource Management can process the following types of Web Log files:

- Common Log Format (CLF) that can be produced by Microsoft and Apache Web Servers
- Extended Log Format (ELF)
- Microsoft IIS original format
- Oracle WebLogic

SAS IT Resource Management examines the file to determine which type it is, without the type having to be specified. However, some Web logs are produced with different delimiters than the expected default. For example, Oracle WebLogic uses a format very similar to Extended Log Format. However, it uses a tab character instead of a space. This delimiter must be specified in the staging transformation properties or supplied as part of the Adapter Setup wizard prompts.

In addition, some Web logs do not provide information about the machine that they were running on or what Web site they were associated with. For this reason, you can also specify the machine and sitename as parameters to the staging transformation or Adapter Setup wizard prompts. Web log files can be concatenated to form one input stream, but they must all have the same Web log format.

Accessing Raw Data for Database Adapters

What Are the Database Adapters?

Several adapters extract raw data from databases. The following adapters are database adapters:

- HP Reporter
- MS SCOM
- VMware
- SAP ERP

Preparing a Database Adapter for Staging

The data sources for the HP Reporter, MS SCOM, and VMware adapters add performance data to relational databases as follows:

- For the VMware adapter, the relational database can be an ORACLE or DB2 database on a UNIX or Windows operating system or a Structured Query Language (SQL) database

- For the HP Reporter adapter, the relational database can be an ORACLE database on a UNIX or Windows operating system, a Structured Query Language (SQL) database, or a Microsoft Data Engine (MSDE) on a Windows operating system.
- For the MS SCOM adapter, the relational database must be a Microsoft SQL Server database.
- For the SAP ERP adapter, the raw data must be in SAS data sets.

Note: If an ODBC server definition must be created that points to that database, see the “ODBC Sources” topic in the “Overview of Common Data Sources” chapter in the *SAS Intelligence Platform: Data Administration Guide* at <http://support.sas.com/documentation/onlinedoc/intellplatform/index.html>. Find the section that corresponds to the type of database that contains your data, and follow the instructions for that section. These instructions include the steps required to set up the ODBC driver on your system if necessary, define the database server, and create the library that points to your database. If you want to establish connectivity to a SQL Server Database by using ODBC, follow the instructions for “Establishing Connectivity to an Oracle Database by Using ODBC.” (The descriptions of ODBC connections to Oracle and Microsoft Access databases can be helpful when you connect to similar data sources.)

The data from one of these databases is input directly into the Adapter Setup wizard. You can set up the library that points to the database in SAS IT Resource Management or in SAS Management Console. After the library is defined, it can be specified as the input database to the appropriate staging transformation.

- To use the SAS IT Resource Management client to define the library that points to the database, follow the instructions in the section called Registering SAS Libraries in the “Getting Started” chapter of the *SAS Data Integration Studio: User’s Guide*. (SAS IT Resource Management uses many features, including registering SAS libraries, that are provided by SAS Data Integration Studio.)

Note: You do not need to specify the metadata for all tables that you want to access in the library. When the staging transformation runs, it specifies the metadata for the tables automatically.

- To use SAS Management Console to define the library that points to the database, follow the instructions that are documented in the “Connecting to Common Data Sources” chapter of the *SAS Intelligence Platform: Data Administration Guide*.

Duplicate-Data Checking for Database Adapters

The staging transformations for the SAP ERP, HP Reporter, MS SCOM, and VMware adapters do not provide a parameter for duplicate-data checking. Duplicate-data checking is always enabled for these adapters.

However, if you want to backload data for HP Reporter, MS SCOM, and VMware staged tables that are added to the PFD after the staging job has already run once, you can override the default action of subsetting the incoming data based on the duplicate-data control data sets. To do so, set the macro variables ITRM_LoadFromDate and ITRM_LoadToDate to valid start and end datetime values. Then, this datetime range will be used to subset the data from the database instead of the ranges in the duplicate data control data sets.

```
%let ITRM_loadFromDate=14FEB2010:00:00:00;
%let ITRM_loadToDate=15FEB2010:23:59:00;
```


Note: When these macro variables are used with the VMware adapter, you must specify the values for these macro variables in Coordinated Universal Time (UTC). (UTC time is the same as Greenwich Mean Time (GMT).)

Notes about Accessing HP Reporter Data

HP Reporter, formerly HP OpenView Reporter, captures enterprise systems management measurement data from Windows systems and many UNIX variants such as HP-UX, SunOS Solaris, IBM AIX, Tru 64 UNIX, and Linux.

For information about the performance areas that are collected by HP Reporter, see [“HP Perf Agent and HP Reporter Adapters” on page 474](#).

If the HP Reporter raw data is in a Microsoft Access database and the SAS IT Resource Management server tier is on a 64-bit Windows machine, you need to set up a PC files library. To do so, perform the following steps:

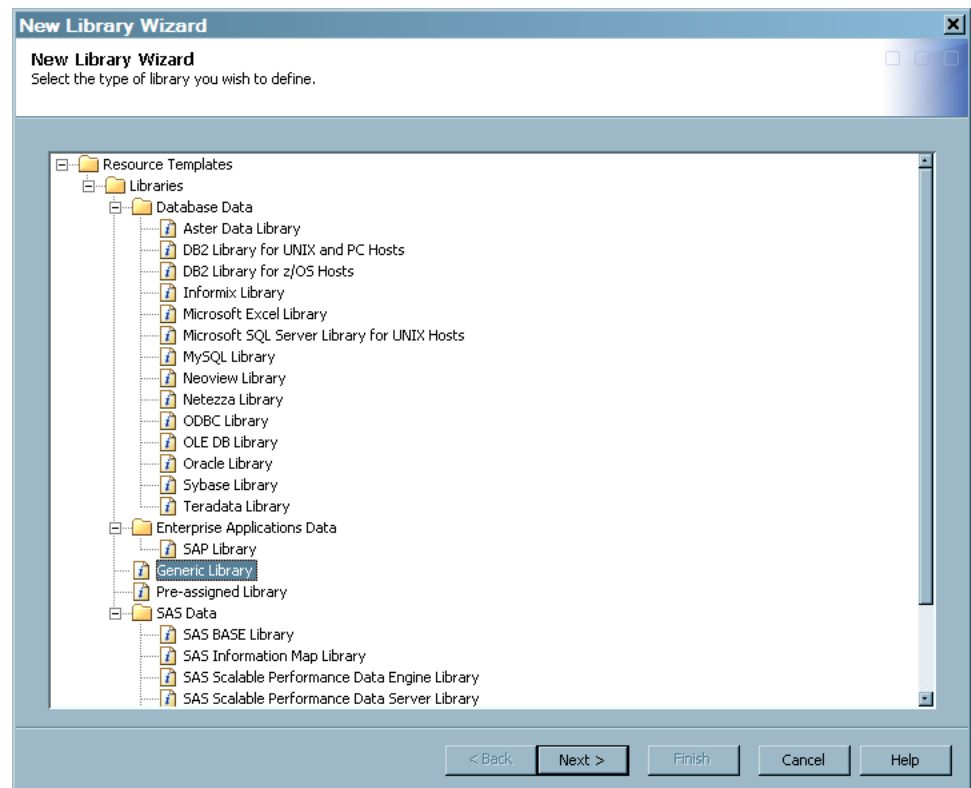
1. Log on to SAS Management Console as an unrestricted user.

To access SAS Management Console, from your Windows desktop, select **Start** ⇒ **Programs** ⇒ **SAS** ⇒ **SAS Management Console**.

2. Navigate to the Data Library Manager.
3. Right-click **Libraries** and select **New Library** from the drop-down list. The New Library Wizard opens.

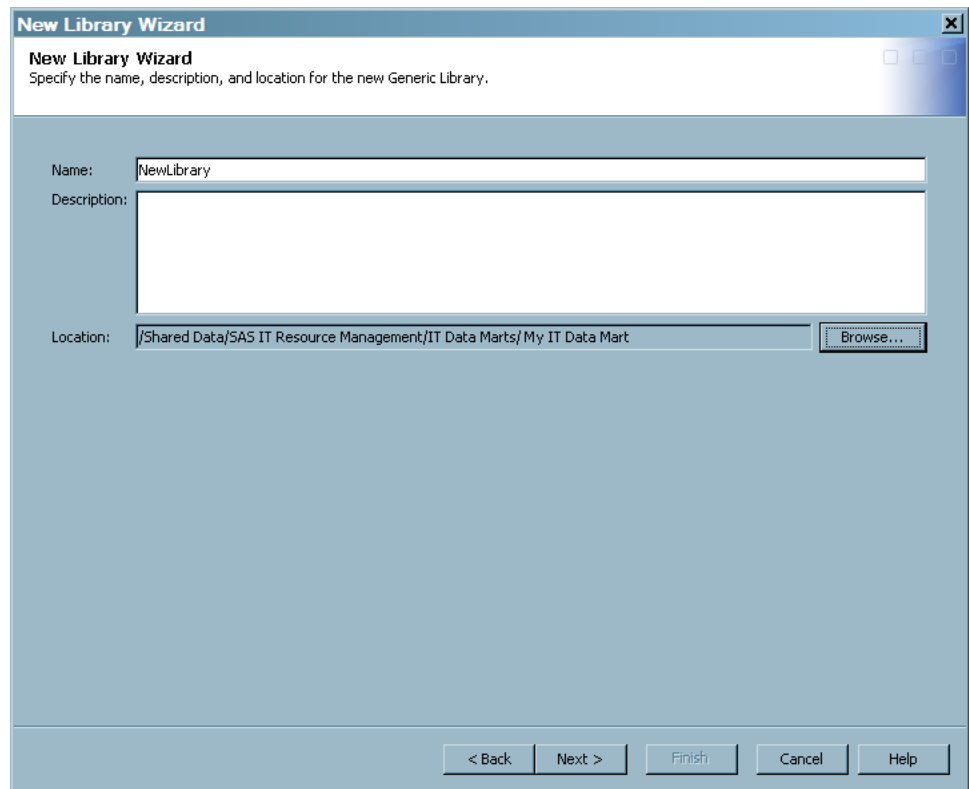
As shown in the following display, select **Generic Library** from the list of Resource Templates.

Display A2.3 New Library Wizard



4. On the next page of the wizard, enter the name of the new library. Use the **Browse** button to display the list of folders from which you can select the **Location** that you want to use for your library.

Display A2.4 Enter Identifying Information for the New Library



New Library Wizard
Specify the name, description, and location for the new Generic Library.

Name:

Description:

Location:

< Back Next > Finish Cancel Help

5. On the next page of the wizard, from the list of available servers select **SASITRM**.
6. On the next page, specify the fields as shown in the following display.

Display A2.5 Enter Information for the Libref

- In the **Libref** field, enter the name by which you want to refer to the new library.
- In the **Engine** field, enter *pcfiles*.
- In the **Option(s) to be appended** field, enter the specifications for the server, port, and path.

The following example shows how to enter these specifications.

```
server="server56.na.co.com" port=9621 path="D:\Public\ReporterDatabase29Feb2008.mdb";
```

Note: You do not need to specify a path in the **Path Specification** box. The path should be entered in the **Option(s) to be appended** field, as shown in the preceding code fragment.

7. Click **Next** and then **Finish** to create the new library.

Notes about Loading Tasks and Events Data for VMware

Limiting the Data in the Events and Tasks Staged Table

TIP Users should take care when staging the VMware Tasks and Events tables due to the potentially huge amount of data that exists in the input database for these tables.

The table in the vCenter database that contains data for the Events staged table contains a record for each event that is the result of a task or alarm in vCenter. The table in the vCenter database that contains data for the Tasks staged table contains a record for each task (power on or off of a virtual machine, alarm created, and so on) that occurs in vCenter. To limit the potential for large amounts of data for these two tables, you can do either one of the following precautions:

- Define a datetime range for which data should be loaded for each of these staged tables.
- Specify the last number of days for which data should be loaded for each of these tables.

Limiting the VMware Data in the Events Table

If the Events staged table is added to a VMware vCenter staging job, do one of the following two procedures:

- To define a datetime range for which data should be loaded for each of these staged tables, edit the generated code for the staging job by adding the following code:

```
%let ITRM_LoadEventsFromDate = <start-datetime>;
%let ITRM_LoadEventsToDate = <end-datetime>;
```

For example,

```
%let ITRM_LoadEventsFromDate = 12APR2010:00:00;
%let ITRM_LoadEventsToDate = 13APR2010:00:00;
```

By setting these two macro variables to a valid datetime range, the staging job would load data for the Events table that falls within the specified range.

- To specify the last number of days for which data should be loaded for each of these tables, edit the generated code for the staging job by adding the following code:

```
%let ITRM_LoadEventsForDays = <number-of-days>;
```

For example,

```
%let ITRM_LoadEventsForDays = 3;
```

By setting this macro variable to a valid positive integer, the staging job loads data for the Events table for the last number of days that you specified. In the preceding example, the staging job populates the Events staged table with data for the last three days.

Note: Specify the values for these macro variables in Coordinated Universal Time (UTC). (UTC time is the same as Greenwich Mean Time (GMT).)

Limiting the VMware Data in the Tasks Table

If the Tasks staged table is added to a VMware vCenter staging job, do one of the two following procedures:

- To define a datetime range for which data should be loaded for each of these staged tables, edit the generated code for the staging job by adding the following code:

```
%let ITRM_LoadTasksFromDate = <start-datetime>;
%let ITRM_LoadTasksToDate = <end-datetime>;
```

For example,

```
%let ITRM_LoadTasksFromDate = 12APR2010:00:00;
%let ITRM_LoadTasksToDate = 13APR2010:00:00;
```

By setting these two macro variables to a valid datetime range, the staging job would load data for the Tasks table that falls within the specified range.

- To specify the last number of days for which data should be loaded for each of these tables, edit the generated code for the staging job by adding the following code:

```
%let ITRM_LoadTasksForDays = <number-of-days>;
```

For example,

```
%let ITRM_LoadTasksForDays = 3;
```

By setting this macro variable to a valid positive integer, the staging job loads data for the Tasks table for the last number of days that you specify. In the preceding example, the staging job populates the Tasks staged table with data for the last three days.

Note: Specify the values for these macro variables in Coordinated Universal Time (UTC). (UTC time is the same as Greenwich Mean Time (GMT).)

Working with the VMware Lookup Table

How to Set Up the VMware Lookup Table

The VMware adapter provides a SAS data table that enables you to provide memory specifications for the hosts in your environment. This table is located by default in the SASHELP library, and contains no data. If you want to use this table to provide memory information, perform the following steps:

1. Locate, or create, the IT data mart that you want to use for staging VMware data.
2. Locate the **Administrative** folder. Right-click the library within that folder. (The library is called **Admin nnnnn**, where *nnnnn* is a random number that was generated when the IT data mart was created.)

Select **View Libname**.

3. Select the entire LIBNAME statement in the displayed box. Then right-click the LIBNAME statement and select **Copy**.
4. Launch an interactive session of SAS. Use the **Paste** function to paste the LIBNAME statement into an Editor window.
5. Submit the LIBNAME statement.
6. In the Explorer window, locate the SAS data table called ITMS_VMware_HostMemory by drilling down into the SASHELP library.
7. Use the **Copy** and **Paste** functions to copy this table to the **Admin nnnnn** library.
8. Now you can edit the table that you copied to the **Admin nnnnn** library. To do so, use VIEWTABLE to add and update rows. (To access VIEWTABLE, double-click the table in the Explorer window.)

The ITMS_VMware_HostMemory lookup table has the following six columns:

- **Cluster**
name of the cluster.
- **Host**
name of the host.
- **EsxHostTotalRam**
the total amount of physical RAM in the ESX host machine that is seen by VMware ESX.
- **EsxHostSystemRam**

the total amount of RAM that is allocated to the VMkernel. This amount does not include virtualization overhead. This amount is the fixed amount of RAM that is available based on the host's total RAM.

- **EsxHostVirtualMachinesRam**

the total amount of RAM that is configured for the Service Console.

- **EsxHostServiceConsoleRam**

the amount of physical memory that is remaining for virtualization. This value is calculated by subtracting the Service Console memory and the System Memory from the total amount of RAM in the host.


The value in the cluster column should be the name of the cluster and the value in the host column should be the name of the host for this observation. (Case does matter. The names in the **Host** column should match the names in vCenter up to the first period. For example, Machine.abcdefg.company.com should be listed as Machine.) You can also update this table programmatically if it has the correct columns, as described previously.

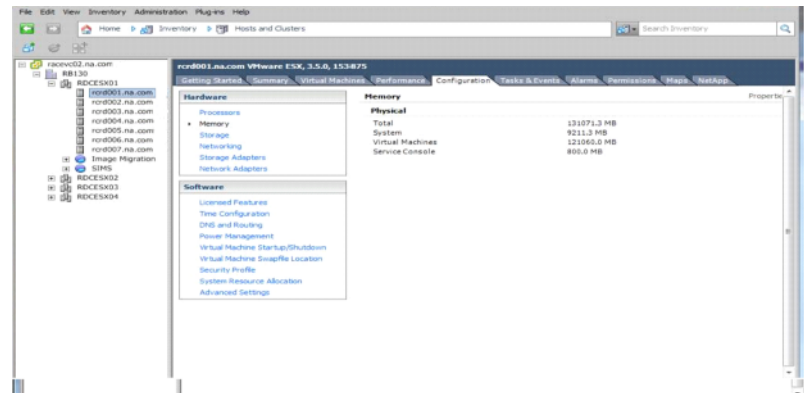
When this table is populated and you run the staging code, the columns **EsxHostTotalRam**, **EsxHostSystemRam**, **EsxHostVirtualMachinesRam**, and **EsxHostServiceConsoleRam** in the HostSystem staged table will have with the correct values from the lookup table. If values for the memory metrics are not provided for a cluster and host, then those variables are set to missing values for that cluster and host. If the lookup table does not exist in the ADMIN library or is empty, then all values for the four mentioned memory variables are set to missing.

Note: If the lookup table does not exist in the ADMIN library or is empty, the HA and DRS supplied reports that depend on the **EsxHostTotalRam**, **EsxHostSystemRam**, **EsxHostVirtualMachinesRam**, and **EsxHostServiceConsoleRam** metrics will not provide useful reports.

How to Populate the VMware Lookup Table

The Virtual Center (vSphere Client) contains information that can be used to load the VMware lookup table. To populate the VMware lookup table, perform the following steps:

1. Log on to the VMware vSphere Client.
2. In the menu bar, select **Inventory** ⇒ **Host & Clusters**. The name of each cluster is displayed in the left panel next to the symbol .
3. For each cluster that is listed, perform steps 4 and 5.
4. Click the + to expand the contents of the cluster. All the hosts that are associated with that cluster are listed.
5. For each host that is listed, perform the following steps:
 - a. Click the name of the host. The Host Details page opens in the right pane. In the right pane, click the **Configuration** tab.

Display A2.6 Host Details Page of VMware vSphere Client

- b. In the **Hardware** section, click **Memory**. The four memory values are displayed. Enter the four values for the appropriate host in the VMware lookup table as follows:
- Copy and paste the **Total** value into the **EsxHostTotalRam** column.
 - Copy and paste the **System** value into the **EsxHostSystemRam** column.
 - Copy and paste the **Virtual Machines** value into the **EsxHostVirtualMachinesRam** column.
 - Copy and paste the **Service Console** value into the **EsxHostServiceConsoleRam** column.

Accessing Raw Data for SAP ERP**Preparing SAP ERP Data for SAS IT Resource Management**

Enterprise Resource Planning (ERP) Software from SAP (formerly SAP R/3), via the SAS IT Management Adapter for SAP, collects SAP ERP and SAP Business Warehouse (BW) workload performance measurements. The SAS server that is used to run the batch ETL work for SAS IT Resource Management can also run the Remote Function Call (RFC) server. The SAP Advanced Business Application Programming (ABAP) code that is executed on the SAP system sends the requested data to the RFC server. The data is stored on that RFC server or on a disk drive that can be accessed on the network.

To enable this communications process, perform the tasks that are listed in the following topics.

Install the Components of ITM Adapter for SAP ERP

Part of the server installation includes transport files that must be moved to the SAP system and installed there. These files include the ABAP programs that enable the extraction of the performance data. For information about how to install these files, see the *SAS IT Management Adapter for SAP: User's Guide*.

The SAS server installation process also installs the `sasrfc_server` and `sasrfc_serveru` executable programs into the SAS tree at `!sasroot\access\sasexe`. One or both of these programs should be installed as Windows services in order to create the RFC server. The `sasrfc_serveru` executable is to be used when the SAP ERP system is running in Unicode.

The SAS client installation includes the `itmadaptsapc` component, which contains the `itmadaptsap.spk` file. When this file is imported into SAS Data Integration Studio, it

contains all the sample jobs that are required to set up and extract the SAP performance data. The client installation also includes plug-ins for SAS Data Integration Studio and SAS Management Console.

Establish Communications between SAS and SAP

To establish communications between SAS and SAP, perform the following steps:

1. Set up an RFC connection on the SAP system.

The SAP administrator should use transaction SM59 to create an RFC Connection to enable communication to your RFC server. Select a TCP connection and enter a meaningful description. For ease of identification, select "Registered Server Program" with a program ID that matches the RFC name. If the SAP system is running in Unicode, ensure that the Unicode flag is set.

2. Create a program variant on the SAP system.

The SAP administrator should use transaction SA38 to create a variant for program / SAS/Z_SAS_READ and update the parameters that are passed to the ABAP programs.

- To facilitate identification, name the program variant after the RFC server.
- For the **G_BUFMAX** parameter, enter 100000.
- For the **G_DEST** parameter, enter the RFC server name.

3. Configure Windows services to create the RFC server. To start the RFC server, use the following type of command: **sasrfc_server -i2 -V <variant-name> -n "<name>,<description>" -p <port_number>** where

- *variant_name* is the name of the variant of the SA38 transaction that you assigned in step 2.
- *name* is a name that you can choose for the service.
- *description* is a description that you can choose for the service. This description field displays as the name of the service when you display the list of services on the Services window. You can access services by selecting **Windows ⇒ Control Panel ⇒ Services**.
- *port_number* is the port number that you have assigned. The default port number is 6999.

Note: If the SAP system is running in Unicode, change the command so that it refers to the Unicode version instead:

```
sasrfc_serveru -i2 -V <variant-name> -n "<name>,<description>" -p <port_number>
```

If a mixture of Unicode and non-Unicode servers is running, install both services, but assign them to different ports.

4. Test the connectivity between your SAS session and the SAP system by using the RFC server that you started in step 3. Start SAS and submit the following code:

```
libname mysap r3 user=<sap_userid> password=<sap_password> client=800
      ashost=<sap_application_server_hostname> sysnr=<sap_system_number>
      host=<rfc_server_host> port=<port_number> batch=Y;
data work.t000;
set mysap.t000;
run;
libname mysap clear;
```


where

- *sap_userid* and *sap_password* are the user ID and password, respectively, that are associated with the SAP system.
- *sap_application_server_hostname* is the host name of the SAP system.
- *sap_system_number* is the system number of the SAP system.
- *rfc_server_host* is the host name of the server that is running the RFC service.
Note: This value is typically *localhost*.
- *port_number* is the assigned port number to be used.

Create SAS Metadata Objects Using SAS Management Console

To create SAS metadata server objects that represent SAP systems, authentication domains, and data libraries, perform the following steps:

1. Invoke SAS Management Console.
2. From the **File** menu, select **Connection Profile**. Then connect to the profile that is being used for SAS IT Resource Management. In the **Repository** field, select **Foundation**.
3. Add an SAP authentication domain for the user ID that you want to be used to access the SAP system. To do this, open the **User Manager**, right-click the user ID and select **Properties**. On the **Accounts** tab, click **New** to add a new domain for SAP authentication. If you have multiple SAP systems with different user IDs or passwords, add one authentication domain per unique user ID and password pair.
4. Define new SAP servers, one for each SAP system. To do this task, perform the following steps:
 - a. Right-click the **Server Manager**.
 - b. Select **New Server** to invoke the New Server wizard.
 - c. From the list of **Resource Templates**, select **SAP Server**.
 - d. On the next page, enter a name for the server and a description. The description is optional.
 - e. On the page that displays server properties, accept the default values.
 - f. On the page that displays connection properties, enter the requested information in the following fields:
 - **Authentication Domain:** select the appropriate authentication domain that was specified in step 3.
 - **Client:** enter the appropriate client. (This value is typically *800*.)
 - **Language:** enter the appropriate language. (This value is typically *EN*.)
 - g. In the **Select Connection Type** part of the page, click the **Application Server** radio button, and select **Options**. The Application Server Options dialog box appears.
 - h. In the **Application Server Host** field, enter the name of the SAP application server host.
 - i. In the **System Number** field, enter the system number. Click **OK**.
 - j. Select **Advanced Options**. In the Advanced Options dialog box, select **RFC Server Batch Mode**. Then click **OK**. Click **Next** to display the information that you specified for this server.

Note: Repeat step 4 in this task for each SAP system. If you have different pairs of user IDs and passwords for the SAP systems, repeat step 3 to ensure that you are using the correct authentication domain for each system.

5. Define new SAP libraries that represent the newly defined SAP servers. To do this task, perform the following steps:
 - a. In SAS Management Console, expand the entry for the **Data Library Manager**.
 - b. Right-click **Libraries**.
 - c. Select **New Library** to invoke the New Library wizard.
 - d. Select **SAP Library** and click **Next**.
 - e. Enter the name and a description for the new library. The description is optional.
 - f. On the Select SAS Server page, select the SAS application server where this library is to be assigned. (This server is typically SASITRM).
 - g. On the next page, enter a unique libref and click **Next**.
 - h. On the next page, select a database server.

(If you want to define a new server, click **New** and follow the instructions in step 4.

Click **Next** to display the specifications for your library.
 - i. Click **Finish**.
6. Create two Base SAS Engine libraries to hold the extracted raw data and other administrative data. To create the libraries, perform the following steps:
 - a. In SAS Management Console, expand the **Data Library Manager** node.
 - b. Right-click **Libraries**.
 - c. Select **New Library** to invoke the New Library wizard.
 - d. Select **SAS Base Library**.
 - e. Enter the name and a description for the new library. The description is optional.
 - f. On the Select SAS Server page, select the SAS application server where this library is to be assigned. (The server that is typically assigned is SASITRM.)
 - g. On the next page, enter a unique libref, such as rawdata or sapadmin.
 - h. In the **Path Specification** part of the page, select a path for the data library. Alternatively, click **New** to create a new data library.
 - i. Click **Next** to display the specifications for the library that you defined.
 - j. Click **Finish** to create the library that you defined.

Configure Sample Data Extraction Jobs Using SAS Data Integration Studio

To configure sample jobs that extract the SAP ERP data, perform the following steps:

1. Start SAS IT Resource Management.
2. Select **File** ⇒ **Import** ⇒ **Metadata** to invoke the Import Metadata Wizard. (You might be prompted to enter the location of your SAS License file.) When prompted by the wizard, enter the following information:
 - Enter the name of the SAS package file to be imported: *itmadapsap.spk*.

- Enter the location where the SAS package was installed: **c:\program files\SAS\SASITMAdapterSAP\3.21\ETLComponentsPackage**. (This location is the default location for the installation of the SAS package file.)
- Clear the **Include access controls** check box.
- Select the option to import **All Objects**.
- Select the SAS application server where this library is to be assigned. (The server that is typically assigned is SASMain.)
- In the **Library** field, enter the SAP library that was defined in the previous step.
- In the fields for the physical locations of the administrative data and the raw data, enter the locations that were created in step 6 of the task called “Create SAS Metadata Objects Using SAS Management Console.”

Select **Finish** to create a folder called **ITM Adapter 3.1 for SAP**.

3. Each sample job is used to achieve a particular step. Samples are provided for scenarios with single SAP servers and multiple SAP servers, so be sure to select the appropriate sample for your environment. Copy and paste only the sample jobs that you need. These sample jobs were installed into the **Jobs** folder in the IT data mart that is designated for the SAP adapter when the Metadata Import Wizard ran in step 2.

Note: The following two SAS options should be added to the SAS config file or to the individual adapter batch jobs when they are executed on a z/OS workspace server:

```
options nocardimage;                      options linesize= 300 ;
```

The following scenario describes the jobs that handle multiple SAP servers:

- Use **0010B Step 1 Select SAP Servers** to identify multiple SAP servers from which you want to extract performance data. Using the **Custom View**, locate the **Shared Data** folder in the Foundation repository. Then, locate the SAP Server Library that you defined in SAS Management Console. Drag that library onto the drop zone labeled “Place library here” in the Process Editor window. Repeat for as many libraries as you defined.
- Use the job named **0020 Step 2 Suggest Destinations for Selected SAP Servers** to extract host information from those SAP servers.
- Use the job named **0030 Step 3 Create Final List of Servers and Destinations** to check for duplicate servers and release compatibility.
- Use the job named **0040 Step 4 Create Mapping for SAP Program types** to create a mapping table of SAP programs to applications.
- Use the job named **0110A Extract Default SAP Performance Data** to actually extract the SAP performance data.

Note: This job is the only one that you should have to deploy or schedule. It extracts the SAP ERP performance data and stores it in the rawdata SAS library. Make sure it is run on a regular basis, often enough to ensure that you lose no performance data. (Some sites run it as frequently as every hour.) This job writes to the rawdata library, appending data each time it is run, so that no data is overwritten.

4. The rawdata library must be specified as the input rawdata library to the SAP ERP staging transformation. The ETL job that this transformation generates reads the rawdata library into the IT data mart. However, it does not clear the contents of the rawdata library.

Recommendation: Automate a backup job that copies the rawdata library to an archive location and then empty the rawdata library before the execution of the next extraction job.

Note: For more detailed information about these tasks, see the *SAS IT Management Adapter for SAP: User's Guide*.

For information about the performance areas that are collected by SAP ERP, see “SAPR3 Adapter - Systems, Machines, Tasks, and Transactions” on page 475 and “SAPR3 Adapter - SAPBTCH (Batch Job) and SAPWKLD (Workload)” on page 477.

Working with Custom Tables

IBM SMF Custom Tables

SAS IT Resource Management provides support for over 1,100 tables in the IBM SMF adapter. However, it does not support every table that MXG can produce from IBM SMF data. For those instances, you can extend the functionality of the IBM SMF staging transformation to load data into user-defined tables for which SAS IT Resource Management does not provide a template. These are called IBM SMF custom tables.

To create an IBM SMF custom table, you must create (or register) metadata about the physical MXG data set or table. You can use the New Table wizard to do so. The physical table must exist first and it must contain a column (variable) called DATETIME to enable it to be added to a user-defined table. If the default definition for the MXG data set does not contain DATETIME, it can be added simply by using MXG exits. This process is documented in the first step of the following task. Once the MXG data set is registered, it can be added as an output table to the IBM SMF staging transformation. It can also subsequently be used as input to an Aggregation transformation.

How to Load IBM SMF Data into User-Defined Tables

To load IBM SMF data into user-defined tables for which SAS IT Resource Management does not provide a template, perform the following steps:

1. If your MXG data set (with the DATETIME variable present) already exists as a physical table, skip this first step.

Otherwise, create it now using regular MXG coding techniques. If the MXG data set does not have a DATETIME variable defined for it, you must identify the variable that contains the datetime stamp for each event or interval that the data represents. Then, modify the EX exit member that contains the OUTPUT statement for that table to also contain an assignment statement of the form:

```
DATETIME=mxg_variable;
```

You must also add the DATETIME variable to the output data set's keep list. To do so, add it to the _Kdddddd. For example:

```
%MACRO _KTY1415 DATETIME
%
```

The preceding example adds the DATETIME variable to the keep list for the TYPE1415 table. Look up the correct names for the EX member and the _Kdddddd macro for your data source in the MXG documentation.

2. Because you do not want the SAS IT Resource Management staging transformation to write to this copy of the MXG data set, make a copy of the physical MXG data set. You can put the copy in any writable library, such as the **Staged** library of your IT data mart.
3. Modify the MXG_PROPS_SMF format that maps the name of the tables to the MXG code that has to be executed. Locate the source for this format in your SASHELP.ITRM.FORMATS catalog. In the format, the value to be formatted is the MXG data set name, and the resulting label is a string of MXG code members and macro names. Add a new line to the format containing your MXG data set name and the name of the MXG code and macros using what is there as a template. Make sure that the text strings start in the same columns. When you execute PROC FORMAT, revise the LIBRARY= parameter in the PROC statement to ensure that the output catalog in the PROC FORMAT statement is pointing to ADMIN.ITRM. This revision of the LIBRARY= parameter creates the format in your ADMIN library that precedes the SASHELP.ITRM.FORMATS in your FMTSEARCH SAS option. Therefore, this format will be used instead of the one supplied at installation time.

The task of allocating the ADMIN library requires you to have Update access to it. The allocation can be achieved only when no other SAS IT Resource Management ETL jobs are running against this IT data mart. You must execute the following statement in a SAS session:

```
libname admin "path_of_ADMIN_library";
```

where "path_of_ADMIN_library" can be found in your SAS IT Resource Management or SAS logs, or by querying the ADMIN library of your IT data mart within SAS IT Resource Management.

4. Start SAS Data Integration Studio and use the Register Tables wizard from the **File** menu to register your MXG data set as a user-defined table.

To register your MXG data set as a user-defined table, perform the following steps:

- a. On the first page of the Register Tables wizard, select **SAS** as the source type.
- b. In the **Library** field, select the library that contains your MXG data set. If you copied your MXG data set to a library that has already been defined in metadata, select the appropriate library from the SAS Library list. If you copied it to a library that has not yet been defined in metadata, and the library is not in the list, then you must define a new SAS Library in metadata now by selecting **New**.
- c. On the following screen, select the MXG data sets that you want to import metadata from.
- d. Choose a folder to store the metadata. Select the **Other Tables** folder in the appropriate IT data mart. It should be the same IT data mart that contains the job to which you want to add this table.
- e. Select **Finish** on the confirmation screen to import the metadata. The MXG data sets should be listed within the **Other Tables** folder of your IT data mart in the SAS Data Integration Studio **Custom** tab.

Note: The table has a slightly different icon associated with it because it is not backed by a SAS IT Resource Management staged table template.

5. Add the new user-defined table as an extra output of the IBM SMF staging transformation in your existing jobs. To do so, perform the following steps:
 - a. Double-click the appropriate job. It appears on the **Diagram** tab of the Job Editor window.

- b. In the Transformations tree of SAS IT Resource Management, expand the **ITRM Transformations** folder and locate the IBM SMF staging transformation. Drag and drop that staging transformation onto the Job Editor window. The staging transformation appears in the process flow diagram for the job.
 - c. Right-click the IBM SMF staging transformation and select **Add output**.
 - d. Drag the new user-defined table from the **Custom** tab to the drop zone that now appears on the process flow diagram.
 - e. Save the job.
6. Run the job or deploy it for subsequent batch execution, as you would any other SAS IT Resource Management job.

Working with MXG Code to Stage Data

About MXG Adapters

Raw data sources that are supported through MXG, such as z/OS data sources, might require that you customize some of the MXG code. These changes might be required in order to stage the data appropriately for the MXG adapters that SAS IT Resource Management provides. SAS IT Resource Management supports the following MXG adapters:

- IBM DCOLLECT
- IBM EREP
- BMC Mainview IMS
- DT Perf Sentry with MXG
- IBM SMF
- ASG TMON2CIC
- ASG TMONDB2
- IBM TPF
- IBM VMMON

MXG Views

MXG views are SQL views that are created automatically by the staging transformations for MXG adapters.

MXG views are also created by the Aggregation transformation for simple aggregations that are based on staged tables from the same adapters in the preceding list. One MXG view is created for each of these staged tables and simple aggregation tables. The MXG views are stored in the same SAS library in which their dependent SAS tables are stored.

Each MXG view shows a SAS IT Resource Management table and the variables for that table in a structure that is equivalent to an MXG data set. Thus, these views enable experienced MXG programmers to use SAS IT Resource Management tables in MXG reporting jobs as if the tables had been created by MXG.

For example, the staged table STAGE.XTY70 is accompanied by an MXG view called STAGE.TYPE70. In addition, if a simple aggregation is built from a staged table, then an MXG view called SIMPLE.TYPE70 is created for the SIMPLE.XTY70 table.

IMACWORK and RMFINTRV Customizations Relevant to the RMF Domain Category of the SMF Adapter

The IMACWORK and RMFINTRV members define the service classes (for those running in goal mode) or performance groups (for those not running in goal mode) that make up the workloads at your site. As with all MXG customizations, copy the member that you want to modify from MXG.SRCLIB to MXGUSER.SRCLIB. Edit these copies according to the instructions or comments therein. Both members contain important notes about the mixing of service classes with reporting classes (in goal mode) and control performance groups with report performance groups (not in goal mode). A good job of customizing these members ensures useful information in the XRMFWKL, XRMFWKP, and XRMFINT tables.

Notice that the RMFINTRV member permits the specification of more workloads than the IMACWORK member permits by itself. Depending on your requirements, you could edit just one, or both, of these members.

If you want to simultaneously process SMF data from multiple machines and if these machines have defined different RMF intervals, then you must also review the setting of the **INTERVAL=** macro parameter of the %VMXGRMFI macro in the RMFINTRV member. The value that you choose for this parameter must be an interval into which all RMF intervals that are being processed divide evenly, without a remainder. For example, if the RMF intervals of two machines are 10 minutes and 15 minutes, use **INTERVAL=HALFHOURL**. Review the comments in the RMFINTRV member for all possible values for **INTERVAL=**. Even if you choose to process the data from different machines in different executions of the same staging job, subsequent analysis of the data is more effective if both sets of the data are processed with the same value for the **INTERVAL=** parameter.

IMACSHFT Customizations Relevant to All Domain Categories of the SMF Adapter

The IMACSHFT member is used to define your shift patterns. However, the SHIFT variable in the IT Resource Management staged tables is populated from a formula that is supplied with IT Resource Management, and is not propagated from MXG. The IT Resource Management shift formula's definition can be changed to any valid SAS expression, or set of DATASTEP statements.

Customizations Relevant to the Jobs Domain Category of the SMF Adapter

IMACSPIN Customizations

The IMACSPIN member is used to specify the number of days for which incomplete jobs are retained in SPIN data sets. They are retained until either the jobs become complete or the time specified in the IMACSPIN member elapses. The default value is zero.

For information, see the comments in the IMACSPIN member.

IMACINTV Customizations

To enable the collection of SMF interval data, you must first ensure that your system is configured to generate the SMF interval accounting records. Work with your systems programming team to determine the changes that might be required in MVS system data set, SYS1.PARMLIB(SMFPRMxx).

When that change is made, you must also modify the IMACINTV MXG member. The default is set so that no TYPE30_V (SMF interval accounting) records are written to SAS data sets. You must copy IMACINTV from MXG.SRCLIB to MXGUSER.SRCLIB and remove the comments from around the OUTPUT statement, as described in the comments in the member.

Note: You can be selective about how many observations (or records) are written, by coding IF and THEN statements. The comments in the member contain examples.

Customizations Relevant to the CICS and DB2 Domain Categories of the SMF Adapter

IMACEXCL and UTILEXCL Customizations

The IMACEXCL member contains the original methods by which modifications to the CICS SMF type 110 record were supported in the MXG code that read them. These original methods have been superseded by the UTILEXCL member. The UTILEXCL member is the preferred method to dynamically create a new IMACEXCL member that is customized to support the actual SMF data being used at your site. The code is designed to run once, to generate the IMACEXCL member. You must then store that member in MXGUSER.SRCLIB to enable the SMF records to be decoded correctly into the SAS data sets.

ASUMCICX Customizations

SAS IT Resource Management has added response time counters to the summarized CICS table that MXG generates by default. To implement this feature, copy the ASUMCICX member from the ITRM.CPMISC pds to your MXGUSER.SRCLIB. Failure to do this causes warning messages to be issued about missing variables.

Note: SAS IT Resource Management is adapted to include the new metrics that the ASUMCICX member creates. This MXG member is stored in the SAS IT Resource Management CPMISC PDS on z/OS, or the **misc** folder on UNIX, or the **sasmisc** folder on Windows. Before you use the ASUMCICX member, you must run the UTILEXCL MXG job. This job creates and implements the IMACEXCL member that is customized for your site. Documentation for ASUMCICX and UTILEXCL can be found inside each member.

IMACUOW Customizations

By default, no ASUMUOW observations (or records) are written to the SAS data set. To process unit-of-work data, you must modify the IMACUOW MXG member. Copy it to your MXGUSER.SRCLIB and comment out (or delete) the first definition of the two macros, _NOOBS and _YESOBS. Then remove the comment indicator from the second definition, so that the macros are defined as follows:

```
MACRO _NOOBS    %
MACRO _YESOBS  %
```


In addition, if you have MQ series data that you want to add to the CICS and DB2 unit-of-work process, modify the code as instructed in IMACUOW to remove the comment indicator from the following statement:

```
%LET MXGMQADD=YES;
```

IMACUOWT Customizations

By default, no ASUMUOWT observations (or records) are written to the SAS data set. To process unit-of-work data from ASG TMON for CICS and SMF DB2, you must modify the IMACUOWT MXG member. Copy it to your MXGUSER.SRCLIB and comment out (or delete) the first definition of the two macros, _NOOBS and _YESOBS. Then remove the comment indicator from the second definition, so that the macros are defined as follows:

```
MACRO _NOOBS %  
MACRO _YESOBS %
```

Customizations Relevant to the Jobs Domain Category of the CA TMS Adapter

By default, there is no single DATETIME variable that represents the date and time of each observation. Copy members IMACTMS5 and EXTMSDSN to your MXGUSER.SRCLIB. In the IMACTMS5 member, add the variable DATETIME to both the _KTMSDSN and _KTMSTMS macros by using the following code:

```
MACRO _KTMSDSN DATETIME %  
MACRO _KTMSTMS DATETIME %
```

In the EXTMSDSN member, insert the following line of code in front of the OUTPUT statement:

```
RETAIN DATETIME "&SYSDATE:&SYSTIME"DT;
```

Customizations Relevant to the ASG TMON2CIC Adapter

If you are processing compressed TMON CICS records, you must do the following:

1. Copy the member EXITMON6 to your MXGUSER.SRCLIB, and follow the instructions therein to assemble and install the MXG INFILE exit "TMON". This action permits both compressed and uncompressed records to be processed without intervention.
2. Copy the member IMACTMO2 to your MXGUSER.SRCLIB. Modify it to remove the comment from the three lines of code as described therein. This change causes the INFILE statement to use the TMON exit described in the preceding step.

Alternatively, the same three lines of code could be specified in the IMACKEEP member instead.

Complete the following customizations in MXG in order to process data for TMON/CICS versions 8.0, 9.0, 1.0 and 1.1:

1. EXITMON6 - If you are using compressed TMON data, then tailor and run EXITMON6 by following the instructions that are contained within EXITMON6.

2. IMACMONI - 1) Tailor IMACMONI by following the instructions within the member.

Complete the following customizations in MXG in order to process data for TMON/CICS ESA version 1.3 and later:

1. EXITMON6 - If you are using compressed TMON data, then tailor and run EXITMON6 by following the instructions contained within EXITMON6.
2. IMACTMON - 1) Tailor IMACTMON by following the instructions within the member.

You can also complete the following customizations to process TMON2CIC data with MXG:

1. Run the job <prefix>.BAMISC(TMONEXIT) and put the module in <prefix>.LIBRARY.
2. Update IMACKEEP by following the instructions in the EXITMON6 member of MXG sourclib. Put the following statements in IMACKEEP:

```
MACRO _TMDB;
INFILE TMDBIN TMON STOPOVER LENGTH=LENGTH COL=COL END=END;
%
```

3. Run the staging job.

Specify Spin Libraries for ASG TMON2CIC and IBM SMF

The ASG TMON2CIC and IBM SMF adapters require that you specify a spin library for spin data. The staging transformations for these adapters create and register a separate spin library. Any existing data is moved to the new spin library as part of the installation, update, or migration process.

For the ASG TMON2CIC and IBM SMF adapters, a spin library must be created. You can create this library on the **Spin Library** tab of the staging transformation properties. On this tab, you can specify various parameters of the spin library or accept the default values. The spin library is then created only after you open the **Spin Library** tab and select **OK**. (If you are using the Adapter Setup wizard to specify the ASG TMON2CIC and IBM SMF adapter, the wizard creates the spin library for you.)

To specify a spin library for ASG TMON2CIC or IBM SMF, perform the following steps:

1. From the Properties dialog box of an ASG TMON2CIC or IBM SMF staging transformation, click the **Spin Library** tab.

Display A2.7 Staging Transformation Properties: Spin Library Tab

ASG TMON2CIC Properties

General | Staging parameters | **Spin Library** | Options | Code | Precode and Postcode | Parameters | Notes | Extended Attributes

Location
IT Data Mart: MyIT Data Mart

Spin library specification
Name: TMON2CIC Spin 95727
Description: Contains spin processing data for job "HP Perf Agent Staging".
Server: SAS
Libref: SPIN
Path Specification: C:\HPPerfAgentStaging\spr95727 Browse

Specify or accept default attributes for allocating SAS libraries on z/OS
File system device:
Primary space, in cylinders:
Secondary space, in cylinders:
Number of volumes: Restore Defaults

OK Cancel Help

2. Confirm that the **IT data mart** is correct for the current job. The value in this field cannot be changed.
3. Specify the **Name** and **Description** of the spin library. SAS IT Resource Management generates default values for these parameters that include a standard naming convention for the library name (**<adapter name> + Spin + <unique number>**). For best results, if you modify the library name, use a standard naming convention that implements unique, identifiable names for all spin libraries that you register. A name value is required and a description is optional.
4. Confirm the values for the **Server** and **Libref** fields for the IT data mart and spin library. The values in these fields cannot be changed.
5. Enter the value for the **Path Specification** field for the metadata path that stores the spin library. SAS IT Resource Management generates the default path **<adapter name> + Spin + <unique number>**. You can accept the default path, enter a new path, or click **Browse** to select a path. A value is required in this field.

Note: **Browse** is disabled when the selected application server is running on a machine using the z/OS operating system. When entering a z/OS path manually, you can use either a prefix for a traditional z/OS filesystem path (**(MY.DATAMART)**) or a root directory in the UFS hierarchical file system (**/u/myname/datamart**).

6. If you are creating and registering libraries on z/OS, specify the following values:
 - file system device
 - the number of cylinders to use as the primary and secondary space allocation for the data set

- the number of volumes to use when allocating space

A value is required for these parameters only if you are working with z/OS.

7. Click **OK**. The staging transformation creates the spin library as specified.

Limit IBM SMF Records

You might want to limit the IBM SMF records that are processed into an aggregation table to specific systems. To do so, you can use an aggregation filter. However, the disadvantage of using an aggregation filter is that the staging code is unnecessarily processing system data that will not be aggregated. In addition, this requires that a filter be set on each and every SMF aggregation. It is more efficient to filter data during the staging job processing.

To limit the IBM SMF records in order to process only specific systems in SAS IT Resource Management during a staging job process, add an IF statement in the IMACFILE MXG member. For example, to process only those records from a 'PROD' and 'TEST' system, perform the following steps:

1. Place a copy of the IMACFILE MXG SOURCLIB member into your user MXG SOURCLIB.
2. Modify the user IMACFILE MXG SOURCLIB member with the following SAS statement: **IF SYSTEM IN ('PROD','TEST'); /*Only process records for these systems */**where the systems for which you want to process IBM SMF records are listed in the IN clause.

For more examples or information about coding the IF statement, see the documentation inside the IMACFILE MXG SOURCLIB member.

Note: IMACFILE is a specific exit point taken by all SMF records, immediately after the SMF header is read. This exit point can be used if you want to bypass the processing of a record based on the machine that produced it.

How MSU and MIPS Columns Are Included in Staged Tables for the SMF Adapter

Overview of the MIPS and MSU Automation Process

In SAS IT Resource Management 3.2 and later, additional columns are included that enhance the reporting of resource utilization for various categories of SMF activity, such as systems, LPARs, workloads, workload periods, jobs, and more. The additional columns reflect MSU (millions of service units) and MIPS (millions of instructions per second) based on CPU time expended during the activity. These columns are distinct from other MSU-related columns that are already present in the SMF rawdata. The following topics explain how the population of these new resource utilization columns in several of the SMF performance-based adapters and tables were automated.

Calculation Dependencies for MIPS and MSU Columns

The calculation of MIPS and MSU columns for SMF staged tables depend on the following items:

- the creation of the MXG RMF interval data set. The following columns are extracted directly or derived from this data:
 - ID of the physical machine (MXG column SYSTEM)

- ID of the physical CEC (MXG column CECSER)
- CPU Model ID (MXG column CPCFNAME)
- CPU Model Capacity Type (Full or Sub, value determined from CPCFNAME value)
- Number of CPUs for the CPU model (MXG column NRCPUS)
- Number of IFA specialty engines configured for use (MXG column NRIFAS)
- Number of ZIP specialty engines configured for use (MXG column NRZIPS)
- Total configured CPU engines. This column is the key merge column with the IBM Large System Performance Reference (LSPR) information.
 - If the CPU model is full capacity, this column is calculated as the sum of all configured standard, IFA, and ZIP engines.
 - If the CPU model is sub-capacity, the MXG column NRCPUS value is used.
- Adjusted CPU Model ID:
 - If the CPU model is full capacity, the value of this column is determined by recalculating CPU engine specification to include the quantity of those additionally configured IFA and ZIP engines.
 - If the CPU model is sub-capacity, the MXG column CPCFNAME value is used.
- Software service unit rate constant (MXG column CECSUSEC)
- Hardware service unit rate constant (MXG column SU_SEC)
- CPU Model MSU capacity (MXG column CPCMSU)
- Minimum datetime present in the RMF interval data for the combination of the previously listed columns
- Maximum datetime present in the RMF interval data for the combination of the previously listed columns
- the availability of Vendor-supplied LSPR Resource Chart from IBM, from which the following columns are extracted:
 - CPU Model ID (This column is labeled Processor in the Vendor-Supplied LSPR Resource Chart. It is a key column that is used for merging with RMF interval data.)
 - CPU model available CPUs (This column is labeled #CP in the resource chart.)
 - IBM Processor Capacity Index (Also known as MIPS, this column is labeled PCI in the Vendor-Supplied LSPR Resource Chart.)
 - CPU Model MSU capacity (Also known as MSU, this column is labeled MSU in the Vendor-Supplied LSPR Resource Chart.)
- the normalization of all CPU times used in MSU and MIPS calculations. (Based on current information, these CPU times are normalized in MXG.)

Initialization and Preparation of the MIPS and MSU Columns

SAS IT Resource Management supplies an initial copy of a SAS data set that contains currently available LSPR information, as well as additional columns that enhance the reporting of resource utilization for various categories of SMF activity. For more information about LSPR tables and how to update the values in them, see <http://support.sas.com/kb/40/936.html>.

Note: IBM might update this information with new processors or other information at a later date. Therefore, the LSPR information might need to be updated. If an update is needed, then the following two steps should be used to create an updated copy of the LSPR information.

To initialize and prepare MIPS and MSU columns for inclusion in your staging tables, perform the following steps:

1. Create a Microsoft Excel spreadsheet that contains selected information from published (and currently available) IBM LSPR information. (You must reference the current published location of the IBM LSPR information. At this time, this information is located at <https://www-304.ibm.com/servers/resourceLink/lib03060.nsf/pages/lsprindex?OpenDocument>.)

Copy and paste a selected set of columns from the published LSPR data into a Windows Excel spreadsheet. The resulting workbook should look like the following table and reflect the information that is documented in the topic called “[Calculation Dependencies for MIPS and MSU Columns](#)” on page 336.

Table A2.2 Vendor-Supplied LSPR Resource Chart

Processor (CPU Model ID)	#CP (Number of CPUs Available per CPU Model)	PCI (IBM Processor Capacity Index)	MSU (CPU Model MSU Capacity)
2066-0E1	1	35	7
2066-0A1	1	71	13
2066-0B1	1	102	20
2066-0C1	1	126	25
2066-0X2	2	145	28
2066-001	1	170	32
2066-0A2	2	234	44
2066-002	2	319	60
2066-003	3	460	84
2066-004	4	590	108
2086-110	1	24	4
2086-210	2	47	8
2086-310	3	70	11
2086-410	4	92	15
2086-120	1	43	7
2086-220	2	84	13

Processor (CPU Model ID)	#CP (Number of CPUs Available per CPU Model)	PCI (IBM Processor Capacity Index)	MSU (CPU Model MSU Capacity)
2086-320	3	125	20
2086-420	4	165	26

2. Once created, the Microsoft Excel spreadsheet can be used as input in the following PROC IMPORT and DATA STEP code. This code creates a permanent SAS data set that resides in an IT Resource Management Admin library.

```

/
/* Read the Microsoft Excel workbook that contains extracted information from
IBM LSPR documentation and create a SAS data set from it. */

PROC IMPORT OUT= WORK.LSPRDATA
    DATAFILE= <"location-of-spreadsheet-created-in-step-3">
    DBMS=EXCEL REPLACE;
    SHEET="Sheet1$";
    GETNAMES=NO;
    MIXED=NO;
    SCANTEXT=YES;
    USEDATE=YES;
    SCANTIME=YES;
RUN;

/* Extract selected fields of information from the vendor-supplied
LSPR resource information. */

DATA WORK.LSPRDATA(keep=CPCFMDL LSPRCPUS LSPRPCI LSPRMSU);
    SET WORK.LSPRDATA;
    LENGTH CPCFMDL $ 8;
    IF _n_=1 THEN DELETE;
    CPCFMDL=F1;
    LSPRCPUS=F2;
    LSPRPCI=F3;
    LSPRMSU=F4;
run;

/* Store IBM LSPR resource information in SAS data set in
IT Resource Management Admin library. */

LIBNAME ADMIN <'location-of-ITRM-Admin-library'>;
PROC SORT DATA=WORK.LSPRDATA out=ADMIN.ITMS_SMF_LSPRDATA;
    BY CPCFMDL;
RUN;

```

Staging Methodology for MIPS and MSU Columns

The following steps describe the methodology that SAS IT Resource Management uses to stage RMF interval data.

1. After MXG stages the RMF interval data and before IT Resource Management begins its staging transformation processing, code is executed to extract the

necessary columns from the RMF interval data (MXG data set RMFINTRV). This code captures the distinct resource utilization entries from the RMF interval data that has just been staged. Another code step also extracts the current contents of the existing Resource Utilization table.

2. Using the recently staged RMF interval data, code is executed to determine whether the available CPU models in the data are full capacity models or sub-capacity models. This model capacity information and a calculation of total available CPUs for the full capacity models are included with the rest of the RMF interval data extracted previously.
3. The existing Resource Utilization table is merged with the staged RMF interval data in preparation for merging with the IBM LSPR information.
4. The latest available LSPR information from IBM is read so that it can be included in the Resource utilization table. The latest available LSPR information is available from one of the following two sources:
 - an IT Resource Management supplied SAS data set in a SASHELP library.
 - an updated copy in an IT Resource Management Admin library that was created by the customer. For information about this process, see [“Initialization and Preparation of the MIPS and MSU Columns” on page 337](#).
5. The combined Resource Utilization table information (from the previous resource utilization table data and the new RMF interval data) and the extracted IBM LSPR resource information are merged together by the common column CPCFMDL, which is the adjusted CPU model. This merge creates a new current copy of the Resource Utilization table, which replaces the existing copy. (The staging code created an empty Resource Utilization Table for this purpose.)

If, as a result of this merge, a machine in your RMF data is not listed in your LSPR table, SAS IT Resource Management will halt the staging job and write an ERROR message to the SAS log.

- If the LSPR information and the subsequent calculations that are based on it are important to you, you should update your LSPR table as soon as possible. Do not attempt to rerun the staging job until the updated table has been implemented and stored in your Admin library. For information about updating your LSPR table, see [“Initialization and Preparation of the MIPS and MSU Columns” on page 337](#).
- If you are not interested in the LSPR information, you can bypass the error and continue processing the data. To do so, set the LSPR_ERROR_ON macro variable at the beginning of your staging job. The following code bypasses the error and enables processing to continue:

```
%LET LSPR_ERROR_ON=N;
```

This code sets all subsequent calculations that are based on the LSPR information to missing values. The SAS log will contain notes about this instead of errors.

6. The newly updated Resource Utilization table data can be used to include various parts of this information in staging extracts so that resource utilization computed columns for MSU and MIPS can be calculated. Lookups can be done using System, CEC, or CPU Model columns present in the Resource Utilization table. The following table shows the type of columns and the corresponding information for those columns that is available in the Resource Utilization table.

Note: The RMF CPU Resource Table consists of 17 columns and a variable number of rows. (In the following example, the table contains sample data for two rows.) In order to present all the columns that are available in the table in a legible format, the

table was rotated 90 degrees, with the columns displaying as rows, and the rows displaying as columns.

Table A2.3 RMF CPU Resource Table: Column Names and Two

Name of the Column in the Resource Utilization Table	Sample Values for First Row of the Table	Sample Values for Second Row of the Table
System ID	DEVA	POA2
Cec ID (from RMF)	B77A	3A7A
CPU Model (from RMF)	2094-706	2066-0A2
Cpu Model Capacity Type (extrapolated from CPU Model ID from RMF)	Full	Sub
Standard CPU Engines (from RMF)	6	2
IFA CPU Engines (from RMF)	2	0
ZIP CPU Engines (from RMF)	1	0
Total CPU Engines (calculated if full capacity model, else standard CPUs from RMF)	9	2
CPU Model extrapolated if full capacity model, else CPU model from RMF)	2094-709	2066-0A2
CPU Model CPUs (from LSPR)	9	2
CEC Su/Sec (CEC Service constant from RMF)	19536.019536	6165.7032755
Su/Sec (CPU Conversion constant from RMF)	25559.105431	6341.6567578
CPU Model MSU Capacity (from RMF)	422.00	44.00

Name of the Column in the Resource Utilization Table	Sample Values for First Row of the Table	Sample Values for Second Row of the Table
CPU Model MSU Capacity (from LSPR)	584.00	44.00
IBM Processor Capacity Index (also known as MIPS from LSPR)	4,366.00	234.00
Minimum Datetime (from RMF)	01Mar08:00:00:00	01Mar08:00:00:00
Maximum Datetime (from RMF)	01Mar08:23:45:00	01Mar08:23:30:00
MSU/MIPS Conversion Factor (calculated)	7.48	5.32

Additional Notes about MIPS and MSU Calculations

The following notes pertain to the calculation of the MIPS and MSU columns.

- When dealing with a full capacity model, any additional specialty CPU engines, such as IFA and ZIP, that have been configured for your system must be included in the calculations. These specialty CPU engines are added to the **Total CPU Engines** and this new value is stored in the resource table under a column labeled **Total CPU Engines (calculated if full capacity model, else standard CPUs from RMF)**.

The **CPU Model ID** value is adjusted to reflect the addition of any specialty engines. For example, suppose there are three additional specialty engines on a full capacity model (2094-706): two engines for IFA and one for ZIP. As a result, the **CPU model** value is adjusted from 2094-706 to 2094-709 to account for the three additional specialty engines. The latter model ID specification is used to match up to the IBM LSPR information to retrieve the relevant MSU and MIPS ratings. These ratings are then used to calculate the MSU/MIPS conversion factor. This refined methodology for the full capacity CPU model results in a different MsuMipsConversionFactor. The difference is due to using an adjusted CPU model and the LSPR MSU and MIPS ratings that are associated with this adjusted CPU model.

In our sample, without using an adjusted CPU model, the calculation of the MsuMipsConversionFactor is 7.32. Using an adjusted CPU model for the calculation, the MsuMipsConversionFactor is calculated to be 7.48. This results in higher MIPS amounts for CPU activity on the full capacity models that include additional specialty engines.

- The calculation of MIPS is done in two distinct ways, depending on the IT Resource Management table that is being staged.
 - For systems (IT Resource Management tables XRMFSYS and XRMFINT) and LPARs (IT Resource Management tables XASM70L and XASMCEL), MSU is calculated using the hardware-based active CPU times from the RMF type 70 records and the hardware service adjustment factor (SUSEC). MIPS are then calculated by multiplying the MSU result by the appropriate MsuMipsConversionFactor for the associated processor type.

- For workloads and workload periods (IT Resource Management tables XRMFWKL, XRMFWKP, XTY72GO) and batch jobs, TSO sessions, and OMVS sessions (IT Resource Management table XJOBS), we calculate MSU by using the application-based active CPU times from the RMF type 72 records and the software service adjustment factor (CECSUSEC). MIPS are then calculated by multiplying the MSU result by the appropriate MsuMipsConversionFactor for the associated processor type.

Note: For details about the specific calculations, see “[MSU and MIPS Calculations](#)”.

- If comparisons are performed between system and LPAR MIPS values that are based on RMF type70 hardware-based active CPU times and other MIPS values, such as workload, workload period, jobs, and so on, that are based on RMF type72 application-based active CPU times, note that the system- and LPAR-based MIPS will be larger. The reason that the times are different is because the uncaptured CPU time is not available in the type72 record. The difference in values is typically in the 5-15% range.

MSU and MIPS Calculations

All calculations listed in the following table are calculated from normalized CPU times, as provided by MXG.

Domain Category	Staged Table	Computed Column	Column Label	Calculation
System	XRMFSYS XRMFINT	SYSMSU	SystemCpuMsuUsage	if CPACTTM > 0 then SYSMSU= CPACTTM * SUSEC / 1000000; else SYSMSU=.;
		SYSMIPS	SystemCpuMipsUsage	if MIPSFACT ^= . and CPACTTM > 0 then SYSMIPS=(CPACT TM * SUSEC / 1000000) * MIPSFACT; else SYSMIPS=.;
		IFAMSU	IfaCpuMsuUsage	if IFAACTT > 0 then IFAMSU= IFAACTT * SUSEC / 1000000; else IFAMSU=.;
		IFAMIPS	IfaCpuMipsUsage	if MIPSFACT ^= . and IFAACTT > 0 then IFAMIPS= (IFAACTT * SUSEC / 1000000) * MIPSFACT; else IFAMIPS=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		ZIPMSU	ZipCpuMsuUsage	if ZIPACTT > 0 then ZIPMSU= ZIPACTT * SUSEC / 1000000; else ZIPMSU=.;
		ZIPMIPS	ZipCpuMipsUsage	if MIPSFACT ^= . and ZIPACTT > 0 then ZIPMIPS= (ZIPACTT * SUSEC / 1000000) * MIPSFACT; else ZIPMIPS=.;
		TOTMSU	AllEngCpuMsuUsage	AllEngCpuMsuUsage if CFACTTM > 0 or IFAACTT > 0 or ZIPACTT > 0 then TOTMSU=(sum(CP ACTTM, IFAACTT, ZIPACTT) * SUSEC) / 1000000; else TOTMSU=.;
		TOTMIPS	AllEngCpuMipsUsage	if MIPSFACT ^= . then TOTMIPS=(sum(CP ACTTM, IFAACTT, ZIPACTT) * SUSEC / 1000000) * MIPSFACT; else TOTMIPS=.;
		MSUCAP	SystemMsuCapacity	DURATION * NRCPUS * SUSEC / 1000000
		MIPSCAP	SystemMipsCapacity	if MIPSFACT ^= . then MIPSCAP=FLOOR(MIPSFACT * (DURATION * NRCPUS * SUSEC / 1000000)); else MIPSCAP=.;
Lpar CecLpar	XASM70L XASMCEL	LPRSYMSU	LparCpMsuUsage	if LCPUPDT > 0 then LPRSYMSU= LCPUPDT * SUSEC / 1000000; else LPRSYMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		LPRYSMIPS	LparCpMipsUsage	if MIPSFACT ^= . and LCPUPDT > 0 then LPRYSMIPS= (LCPUPDT * SUSEC / 1000000) * MIPSFACT; else LPRYSMIPS=.;
		LPRIFAMSU	LparIfaMsuUsage	if IFAACTT > 0 then LPRIFAMSU=IFA ACTT * SUSEC / 1000000; else LPRIFAMSU=.;
		LPRIFAMIPS	LparIfaMipsUsage	if MIPSFACT ^= . and IFAACTT > 0 then LPRIFAMIPS=(IFA ACTT* SUSEC / 1000000) * MIPSFACT; else LPRIFAMIPS=.;
		LPRZIPMSU	LparZipMsuUsage	if ZIPACTT > 0 then LPRZIPMSU=ZIPA CTT * SUSEC / 1000000; else LPRZIPMSU=.;
		LPRZIPMIPS	LparZipMipsUsage	if MIPSFACT ^= . and ZIPACTT > 0 then LPRZIPMIPS=(ZIP ACTT * SUSEC / 1000000) * MIPSFACT; else LPRZIPMIPS=.;
		LPRTOTMSU	LparAllEngMsuUsage	if LCPUPDT > 0 or IFAACTT > 0 or ZIPACTT > 0 then LPRTOTMSU=(su m(LCPUPDT, IFAACTT, ZIPACTT)* SUSEC) / 1000000; else LPRTOTMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		LPRTOTMIPS	LparAllEngMipsUsage	if MIPSFACT ^= . then LPRTOTMIPS=(sum(LCPUPDT, IFAACTT, ZIPACTT) * SUSEC / 1000000) * MIPSFACT; else LPRTOTMIPS=.;
Workload	XRMFWKL	WKLMSU	WorkloadCpuMsuUsage	Workload XRMFWKL WKLMSU WorkloadCpuMsuUsage if WKLCPU > 0 then WKLMSU=(WKL CPU*CECSUSE) / 1000000; else WKLMSU=.;
		WKL MIP S	WorkloadCpuMipsUsage	if MIPSFACT ^= . and WKL CPU > 0 then WKL MIP S=((WKL CPU*CECSUSE) / 1000000) * MIPSFACT; else WKL MIP S=.;
		WKLIFAMSU	WorkloadIfaCpuMsuUsage	if WKLIFA > 0 then WKLIFAMSU=(W KLIFA*CECSUSE) / 1000000; else WKLIFAMSU=.;
		WKLIFAMIPS	WorkloadIfaCpuMipsUsage	if MIPSFACT ^= . and WKLIFA > 0 then WKLIFAMIPS=((W KLIFA*CECSUSE) / 1000000) * MIPSFACT; else WKLIFAMIPS=.;
		WKLZIPMSU	WorkloadZipCpuMsuUsage	if WKLZIP > 0 then WKLZIPMSU=(W KLZIP*CECSUSE) / 1000000; else WKLZIPMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		WKLZIPMIPS	WorkloadZipCpuMipsUsage	if MIPSFACT ^= . and WKLZIP > 0 then WKLZIPMIPS=((WKLZIP*CECSUSE) / 1000000) * MIPSFACT; else WKLZIPMIPS=.;
		WKLTOTMSU	WorkloadAllEngCpuMsuUsage	if WKPCPU > 0 or WKPIFA > 0 or WKPZIP > 0 then WKLTOTMSU=(sum(WKPCPU,WKPIFA,WKPZIP) * CECSUSE) / 1000000; else WKLTOTMSU=.;
		WKLTOTMIPS	WorkloadAllEngCpuMipsUsage	if MIPSFACT ^= . then WKLTOTMIPS=(sum(WKPCPU,WKPIFA,WKPZIP) * CECSUSE / 1000000)*MIPSFACT; else WKLTOTMIPS=.;
Workload Period	XRMFWKP	WKPMISU	WorkloadPeriodCpuMsuUsage	if WKPCPU > 0 then WKPMISU=(WKPCPU*CECSUSE) / 1000000; else WKPMISU=.;
		WKPMIPS	WorkloadPeriodCpuMipsUsage	if MIPSFACT ^= . and WKPCPU > 0 then WKPMIPS=((WKPCPU*CECSUSE) / 1000000) * MIPSFACT; else WKPMIPS=.;
		WKPIFAMSU	WorkloadPeriodIfaCpuMsuUsage	if WKPIFA > 0 then WKPIFAMSU=(WKPIFA*CECSUSE) / 1000000; else WKPIFAMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		WKPIFAMIPS	WorkloadPeriodIfa CpuMipsUsage	if MIPSFACT ^= . and WKPIFA > 0 then WKPIFAMIPS=((W KPIFA*CECSUSE) / 1000000) * MIPSFACT; else WKPIFAMIPS=.;
		WKPZIPMSU	WorkloadPeriodZi pCpuMsuUsage	if WKPZIP > 0 then WKPZIPMSU=(WK PZIP*CECSUSE) / 1000000; else WKPZIPMSU=.;
		WKPZIPMIPS	WorkloadPeriodZi pCpuMipsUsage	if MIPSFACT ^= . and WKPZIP > 0 then WKPZIPMIPS=((W KPZIP*CECSUSE) / 1000000) * MIPSFACT; else WKPZIPMIPS=.;
		WKPTOTMSU	WorkloadPeriodAll EngCpuMsuUsage	if WKPCPU > 0 or WKPIFA > 0 or WKPZIP > 0 then WKPTOTMSU=(su m(WKPCPU,WKPI FA,WKPZIP) * CECSUSE) / 1000000; else WKPTOTMSU=.;
		WKPTOTMIPS	WorkloadPeriodAll EngCpuMipsUsage	if MIPSFACT ^= . then WKPTOTMIPS=(su m(WKPCPU,WKPI FA,WKPZIP) * CECSUSE/ 1000000)*MIPSFA CT; else WKPTOTMIPS=.;
Goal Mode Period	XTY72GO	PRDMSU	PeriodCpuMsuUsa ge	if CPUTM > 0 then PRDMSU=CPUTM *CECSUSE/ 1000000; else PRDMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		PRDMIPS	PeriodCpuMipsUsage	if MIPSFACT ^= . and CPUTM > 0 then PRDMIPS=(CPUTM*CECSUSE/ 1000000) * MIPSFACT; else PRDMIPS=.;
		PRDIFAMSU	PeriodIfaCpuMsuUsage	if CPIFATM > 0 then PRDIFAMSU=CPIFATM*CECSUSE/ 1000000; else PRDIFAMSU=.;
		PRDIFAMIPS	PeriodIfaCpuMipsUsage	if MIPSFACT ^= . and CPIFATM > 0 then PRDIFAMIPS=(CPIFATM*CECSUSE/ 1000000) * MIPSFACT; else PRDIFAMIPS=.;
		PRDZIPMSU	PeriodZipCpuMsuUsage	if CPZIPTM > 0 then PRDZIPMSU=CPZIPTM*CECSUSE/ 1000000; else PRDZIPMSU=.;
		PRDZIPMIPS	PeriodZipCpuMipsUsage	if MIPSFACT ^= . and CPZIPTM > 0 then PRDZIPMIPS=(CPZIPTM*CECSUSE/ 1000000) * MIPSFACT; else PRDZIPMIPS=.;
		PRDTOTMSU	AllEngCpuTimeMsuUsage	if CPUTM > 0 or CPIFATM > 0 or CPZIPTM > 0 then PRDTOTMSU=(sum(CPUTM,CPIFATM,CPZIPTM) * CECSUSE) / 1000000; else PRDTOTMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		PRDTOTMIPS	AllEngCpuTimeMipsUsage	if MIPSFACT ^= . then PRDTOTMIPS=(sum(CPUTM,CPIFATM,CPZIPTM) * CECSUSE/1000000)*MIPSFACT; else PRDTOTMIPS=.;
Batch Jobs TSO OMVS	XJOBS	SYMSU	SystemCpuMsuUsage	if CPUTM > 0 then SYMSU=CPUTM *CECSUSE/1000000; else SYMSU=.;
		SYMIPS	SystemCpuMipsUsage	if MIPSFACT ^= . and CPUTM > 0 then SYMIPS=(CPUTM*CECSUSE/1000000) * MIPSFACT; else SYMIPS=.;
		IFAMSU	IfaCpuMsuUsage	if CPIFATM > 0 then IFAMSU=CPIFATM*CECSUSE/1000000; else IFAMSU=.;
		IFAMIPS	IfaCpuMipsUsage	if MIPSFACT ^= . and CPIFATM > 0 then IFAMIPS=(CPIFATM*CECSUSE/1000000) * MIPSFACT; else IFAMIPS=.;
		ZIPMSU	ZipCpuMsuUsage	if CPZIPTM > 0 then ZIPMSU=CPZIPTM *CECSUSE/1000000; else ZIPMSU=.;

Domain Category	Staged Table	Computed Column	Column Label	Calculation
		ZIPMIPS	ZipCpuMipsUsage	if MIPSFACT ^= . and CPZIPTM > 0 then ZIPMIPS=(CPZIPTM*CECSUSE/ 1000000)*MIPSFACT; else ZIPMIPS=.;
		TOTMSU	AllEngCpuMsuUsage	if CPUTM > 0 or CPIFATM > 0 or CPZIPTM > 0 then TOTMSU=(sum(CPUTM,CPIFATM,CPZIPTM) * CECSUSE) / 1000000; else TOTMSU=.;
		TOTMIPS	AllEngCpuMipsUsage	if MIPSFACT ^= . then TOTMIPS=(sum(CPUTM,CPIFATM,CPZIPTM) * CECSUSE/ 1000000)*MIPSFACT; else TOTMIPS=.;

VMware vCenter Data Acquisition Jobs

Overview of VMware Data Acquisition Jobs

Although staging jobs for SAS IT Resource Management adapters are typically run once a day, the VMware vCenter adapter requires the ability to run staging jobs more frequently than once a day. The reason for this frequency is that staged VMware data must include information about the mapping of a host (physical machine) to the guest (virtual machine). As virtual machines are moved from host to host, the VMware database is updated with the current location of the virtual machine. The VMware database stores only the most recent host-to-guest mapping information. It does not store any previous mappings. Therefore, if the host-to-guest mappings change more than once a day, those mappings must also be extracted from the VMware database more than once a day. Those mappings must also be stored, along with some of the corresponding data, in a table for use as source tables in the nightly run of the VMware staging transformation. The VMware vCenter Data Acquisition transformation provides this functionality.

How Does the VMware Data Acquisition Job Work?

Like all other transformations, the VMware vCenter Data Acquisition transformation is contained in a job. When the VMware Data Acquisition job runs, it extracts a subset of VMware data from a database. The target table that is created by executing the VMware vCenter Data Acquisition transformation is input to the VMware vCenter staging transformation. Unlike the staged tables created by other staging transformations, the target table is not purged when the VMware Data Acquisition job is executed. The VMware vCenter Data Acquisition staging transformation prompts for the following information:

- **Raw data input library** includes two fields (Library and Libref) that specify the SAS library and corresponding libref for the appropriate adapter database. You can enter the library path or use the browsing function to locate and select it. This parameter does not have a default value.

Note: For VMware Data Acquisition, this parameter supports only libraries with engine types of ODBC, OLE DB, DB2, and Oracle.

- **Age limit for data acquisition table (in days)** specifies the number of days that data is kept in the data acquisition table that populates the staged tables. The default value is 7.

The target table that is created is a staged table. It should be used as input to the VMware vCenter staging transformation. When this staged table is input to the VMware vCenter staging transformation, it is merged with other VMware tables by the VMware vCenter staging transformation during the nightly staging run for that adapter.

The VMware Data Acquisition staged table has its own library. The Adapter Setup wizard creates the two staging libraries—one for the VMware vCenter Data Acquisition table, and the other library for the other VMware vCenter staged tables. (If created by the Adapter Setup wizard, these tables and libraries reside in the **Staging** folder of the IT data mart. If created manually, outside of the Adapter Setup wizard, these tables and libraries reside in the location of your choice.)

Preserving Information about the VMware Environment

The VMware DataAcquisition staged table that is output from the VMware Data Acquisition staging job contains useful information about the topology of the VMware environment at the time that the VMware Data Acquisition staging job is run. The topology of the VMware environment includes the names and IDs of each guest, and the host or cluster that each guest belongs to. However, the VMware DataAcquisition staged table does not preserve the information about the movements of a guest from host to host or cluster to cluster between staging runs.

If you want to preserve this information over a period of time, you can instruct the VMware Data Acquisition staging job to create a separate data set to contain this information—the VMwareTopologyDimensionTable. To specify the creation of that table, set a macro variable called CreateTopologyDimensionTable to YES in the generated code (or deployed job code) for the VMware Data Acquisition staging job. To specify that the VMwareTopologyDimensionTable be created, add the following line to the generated code for the VMware Data Acquisition staging job:

```
%let CreateTopologyDimensionTable = YES;
```

The table will be stored in the same staging library as the DataAcquisition staged table that is output from the VMware Data Acquisition staging job.

Note: The VMware Data Acquisition staging job should be scheduled to run as often as is necessary to capture the information about the movements of guests between hosts and clusters in the VMware environment.

You can use this table as required at your site. The table is not used by the VMware staging code.

When Should the VMware Data Acquisition Job Run?

The VMware Data Acquisition job should be scheduled to run whenever it is necessary to capture the movements of virtual machines across physical machines in your VMware environment. If your environment changes rapidly, then the VMware Data Acquisition job might need to run as frequently as once an hour. However, if the virtual machines in your environment do not move between hosts as often, then the VMware Data Acquisition might need to run only once per day. This scheduling is dependent on, and should be arranged according to, the needs of your site.

For information about deploying jobs, scheduling jobs to run in batch mode, and redeploying jobs, see [“Jobs That Process IT Data” on page 263](#).

Appendix 3

Staging Parameters

Staging Parameters	356
Age Limit for Data Acquisition Table	356
Allow Duplicate ID Variables	356
Choose Access Command	356
Choose Raw Data Input Type	357
Default Duration	357
Delimiter in Raw Data	357
Duplicate Checking	357
FTP	358
Future Data	359
JES	360
Logical Record Length	360
Machine	360
Presummarization Duration	360
Raw Data Input Directory	360
Raw Data Input File	361
Raw Data Input Library	361
rsh Host Command	361
Site Name	361
Use Intermediate Staging View	361
Use snmpwalk to Gather Character Data	361
User-written Staging Parameters	362
Staging Parameters for Supported Adapters	362
ASG TMONC2CIC Staging Parameters	362
ASG TMONDB2 Staging Parameters	363
BMC Mainview IMS Staging Parameters	364
BMC Perf Mgr Staging Parameters	365
CA TMS Staging Parameters	365
DT Perf Sentry Staging Parameters	366
DT Perf Sentry with MXG Staging Parameters	367
HP Perf Agent Staging Parameters	368
HP Reporter Staging Parameters	368
IBM DCOLLECT Staging Parameters	369
IBM EREP Staging Parameters	370
IBM SMF Staging Parameters	371
IBM TPF Staging Parameters	372
IBM VMMON Staging Parameters	373
MS SCOM Staging Parameters	374
SAP ERP Staging Parameters	374
SAR Staging Parameters	374
SNMP Staging Parameters	375

User-Written Staging Parameters	376
VMware Data Acquisition Parameters	379
VMware vCenter Parameters	379
Web Log Staging Parameters	380

Staging Parameters

The following topics describe the various parameters that adapters use to stage the raw data. The appropriate combination of these staging parameters is available from the Adapter Setup wizard and from the **Staging Parameters** tab of the Properties dialog box for an adapter or staging transformation. If a parameter includes a browse function, the browse function might be disabled in some instances such as when the input data is associated with an application server that resides on either a traditional MVS operating environment or a UFS file system.

Age Limit for Data Acquisition Table

Age limit for data acquisition table (in days) specifies the number of days that data is kept in the data acquisition table that populates the staged tables.

This parameter is relevant only to the VMware Data Acquisition because other staging transformations purge the existing data of their corresponding staged tables and replace it with new data on execution of the staging job. However, the VMware Data Acquisition staging transformation populates staged tables that are not purged on execution of the staging job. Therefore, this parameter enables you to specify an age limit for the data in this table, after which the data is purged when the staging job executes, so that the table does not grow indefinitely.

Allow Duplicate ID Variables

Allow duplicate ID variables specifies whether duplicate ID variables are permitted when transposing data for the adapter. The value for this parameter can be **Y** (yes) or **N** (no). **N** is the default value.

Choose Access Command

Choose access command, relevant only to the SNMP adapter, specifies the type of raw data such as NNM, NetView, or RRDTool. The following two options for this parameter also govern other staging parameters that are available for configuration:

- **HP NNM / Netview**, the default value, indicates that the raw data is from HP NNM or NetView. If you select this value, then the corresponding field **snmpColDump executable** appears and enables you to specify the location path and command for snmpColDump. The **snmpColDump executable** parameter does not have a default.

Note: If you select **HP NNM / Netview**, then you must have HP OpenView installed on the server where the staging job is running either interactively or in batch mode.

- **RRDTool** indicates that the raw data is stored in a Round Robin Database (RRD) that is managed by RRDTool. If you select this value, then the following fields appear:

- **rrdtool executable** appears and enables you to specify the location path and command for RRDTool. The **rrdtool executable** parameter does not have a default.
- **Number of days to load** specifies the number of days of raw data that you want to backload into the staged table. This field accepts only integers from one to 365. The default value is 2.

Choose Raw Data Input Type

Choose raw data input type specifies whether the raw data is a file or directory that is available from the client network or if the raw data input is available using an FTP access method. The following options for this parameter also govern several of the other staging parameters that are available for configuration. The options that are available for selection vary based on the adapter.

- **File** indicates that the raw data input for the adapter is a file. If you select this value, then the corresponding field **Raw data input file or directory** appears and enables you to specify the full pathname of the raw data file or the directory. You can enter the path directly or use the browsing function to locate the path and enter it automatically. This parameter does not have a default value.
- **File or directory** indicates that the raw data input for the adapter is a file or directory. If you select this value, then the corresponding field **Raw data input file or directory** appears and enables you to specify the full pathname of the raw data file or the directory. You can enter the path directly or use the browsing function to locate the path and enter it automatically. This parameter does not have a default value.
- **FTP** indicates that the raw data input for the adapter is available using the FTP access method. For more information about additional parameters that are required if you select the FTP access method, see [“FTP” on page 358](#).

Default Duration

Default duration, relevant only to the HP Reporter adapter, specifies the value (in seconds) for the duration of the intervals if the input data does not contain an INTERVAL variable. If the input data does not contain an INTERVAL variable and a value for the default duration is not specified, then the interval duration defaults to 3600 seconds, or one hour.

This value must be numeric. If specified, the integer value must be greater than or equal to 1 and less than or equal to 2,147,483,647.

Delimiter in Raw Data

Delimiter in raw data specifies the delimiter (**Space**, **Comma**, or **Tab**) that is used in the raw input data. The values that are available for this parameter vary based on the adapter.

Duplicate Checking

Duplicate checking specifies whether to check for duplicate data and indicates what to do when duplicate data is encountered.

Note: Duplicate checking is automatically enabled for the adapters that use a database as a raw data source (such as HP Reporter, SAP ERP, MS SCOM, and VMware). This parameter is not available for configuration with these staging transformations. For more information about how to override the default value for duplicate checking with these adapters, see [“Example 3: Using Macro Variables to Subset Data for HP Reporter, MS SCOM, and VMware Adapters” on page 415](#).

Here are the values that are available for this parameter:

Discard

removes duplicates. This is the default value for most adapters.

Force

loads data regardless of whether duplicates are found.

Inactive

does not check for duplicates.

Terminate

ends the job if duplicates are found.

CAUTION:

Although the job terminates, the resulting staged table might contain data. However, termination of the job indicates an error in processing, and any data in the resulting staged table might be invalid.

For more information about duplicate data checking, see [“Duplicate-Data Checking Overview” on page 391](#).

Note: If you update the source code of a job and modify the mode for duplicate checking, then you can save the functionality of the new source code to the local file system. However, SAS Data Integration Studio preserves the original source code. The new value of the mode for duplicate checking is not updated in the repository, and the new value is not reflected in the user interface.

FTP

Overview

If you select **FTP** for the **Choose raw data input type** parameter, then additional parameters appear on the parameters page to facilitate the FTP process. The following topics describe these additional parameters that are available for configuration:

Host

Host specifies the name of the remote host. This parameter does not have a default value.

Port

Port specifies the FTP port number. The default value is **21**.

User

User specifies the user ID for the FTP server. This field accepts alphanumeric characters. This parameter does not have a default value.

Password

Password specifies the password of the given user ID for the FTP server. This parameter does not have a default.

CAUTION:

The Adapter Setup wizard generates a job log that displays all information that you specify in the wizard, including any passwords for accessing FTP data. If you have concerns about this password showing in the job log, then you can set up a user ID that you use only for accessing the files via FTP.

External File Name

External file name specifies the filename of the raw data. The maximum number of characters in this field (including the dots in a z/OS filename, if applicable) is 44. This parameter does not have a default value.

Tape

Tape specifies whether the data file is on tape. The default value is **No**.

RCMD

RCMD specifies the FTP SITE or service command that is sent to the FTP server to provide services that are system-specific or essential to transfer files but not common enough to be included in the protocol. The default value is **SITE RDW**.

Debug

Debug specifies whether to write to the SAS log any messages that are sent to and from the FTP server. The default value is **No**.

Future Data

Future data specifies whether to check for future data and indicates what to do when it is encountered. This parameter specifies the FUTURE parameter in the RMFUTURE macro.

This parameter controls the processing of incoming data that has a datetime variable that is greater than 48 hours in the future. (That is, the datetime variable is more than 48 hours after the current time on the system where data is being staged for processing into the IT data mart.) The 48-hour buffer provides for different time zones, daylight saving time, Greenwich Mean Time, and so on.

If future data is encountered, a note is written to the SAS log. This note provides the future data option that is selected, shows the datetime that was encountered, and explains the status of the future data, such as whether it was added to the IT data mart or if the job was terminated.

Here are the values that are available for this parameter:

Accept

specifies that incoming data is staged for processing and is processed into the IT data mart. If any of the data has a datetime value of 48 hours or more in the future, then a note that future data was encountered is written to the SAS log. This value enables an IT data mart to accept future data. For example, you might want to use this setting to perform end-of-year testing with a test IT data mart.

Note: Age limits take effect from the most recent data. Therefore, dates in the future might cause at least some of the existing data to be aged out of the IT data mart.

Discard

specifies that data with a datetime value of 48 hours or more in the future is not staged for processing and is not processed into the IT data mart. This value is the default. This value prevents future data from being processed into the IT data mart.

Future data might cause existing data to be aged out (the existing data would appear to be older than it is, in comparison with the future data).

Terminate

specifies that if any incoming data has a datetime value of 48 hours or more in the future, then staging of the data stops, an error message is written to the SAS log, and the job terminates. This value prevents future data from being processed into the IT data mart, which might cause existing data to be aged out (the existing data would appear to be older than it is, in comparison with the future data). This value stops processing and thus calls more attention to the future data than the **Discard** value.

CAUTION:

Although the job terminates, the resulting staged table might contain data. However, termination of the job indicates an error in processing, and any data in the resulting staged table might be invalid.

JES

JES specifies the version (**JES2** or **JES3**) of the z/OS job entry subsystem that is in use by the system where the input data was recorded.

Logical Record Length

Logical record length specifies the logical record length for the input file. If specified, the parameter value must be a positive integer greater than one. This parameter is pertinent to adapters that are z/OS based for raw data sources.

For several adapters, the default value is left blank to indicate that the logical record length is dependent on the type of input file and operating system that is used.

Machine

Machine specifies the name of the machine that generated the raw data. This parameter is primarily relevant to the Web Log adapter. If the Web Log data already contains a value for machine name, then the value in the raw data is used with the staged table and you do not need to specify a value for this parameter. If the Web Log data does not already specify a machine name, then you can use this parameter to specify a machine name and that value is then associated with the staged table. This parameter does not have a default value.

Presummarization Duration

Presummarization duration specifies the duration, in seconds, of the intervals into which you want to summarize the raw data before it is staged. For example, if you enter *3600*, then the raw data is summarized in intervals of one hour. If specified, the value must be a positive integer less than or equal to 86,400, which is the number of seconds in a day.

Raw Data Input Directory

Raw data input directory specifies the full pathname of the directory for the raw data. You can enter the path directly or use the browsing function to locate and select it.

Raw Data Input File

Raw data input file specifies the full pathname of the raw data file for the adapter. You can enter the path directly or use the browsing function to locate and select it.

Raw Data Input Library

Raw data input library includes two fields (**Library** and **Libref**) that specify the SAS library and corresponding libref for the appropriate adapter database. You can enter the library path or use the browsing function to locate and select it.

Note: For the HP Reporter adapter, this parameter supports only libraries with engine types of ODBC, OLE DB, and Oracle. For the SAP ERP adapter, this parameter supports only libraries with engine types of SAS or BASE.

rsh Host Command

rsh host command, relevant only to the SNMP adapter, specifies the **rsh** command and the name of the host for running the **NNM snmpColdDump** command or the **rrdtool** command. The command in this field is used at the beginning of the **rrdtool**, **snmpColdDump**, or **snmpwalk** commands. This field can be left blank if an **rsh** command is not required to run the **rrdtool** or **snmpColdDump** command.

Site Name

Site name, relevant to the Web Log adapter, specifies the name of the Web site that generated the raw data. If the Web Log data already contains a Web site value, then the value in the raw data is used with the staged table and you do not need to specify a value for this parameter. If the Web Log data does not already specify a Web site, then you can use this parameter to specify a Web site value and that value is then associated with the staged table. This parameter does not have a default value.

Use Intermediate Staging View

Use intermediate staging view specifies whether to use the temporary view form when instantiating staged tables. Enter the staged table name or leave this field blank.

This optional parameter can contain the SAS data set name of a staged table or a macro variable that resolves to a SAS data set name. If not specified, this parameter defaults to a blank, which implies that you do not want presummarization.

Use snmpwalk to Gather Character Data

Overview

Use snmpwalk to gather character data specifies whether to use the **snmpwalk** command to capture data during the execution of the job. The following values are available:

- **No** does not use the **snmpwalk** command.

- **Yes** uses the `snmpwalk` command. If you select this value, then additional parameters appear. The following subtopics describe each of these additional parameters.

This parameter and others associated with `snmpwalk` are relevant only to the SNMP adapter.

snmpwalk Executable

The **`snmpwalk executable`** parameter indicates the location and command for `snmpwalk`. This parameter does not have a default value.

HostFile for snmpwalk

HostFile for `snmpwalk` indicates the location of the `snmphost` file that lists the hosts from which you want to get data. This parameter does not have a default.

Community Value for snmpwalk

Community value for `snmpwalk` indicates the community name that is required to run the `snmpwalk` command. The default value is `public`.

User-written Staging Parameters

For a detailed description of the staging parameters for the User-written Staging transformation, see “[Staging Parameters Tab](#)” on page 234.

Staging Parameters for Supported Adapters

The following tables show the staging parameters for each adapter that SAS IT Resource Management supports. These staging parameters are available from the Adapter Setup wizard and from the **Staging Parameters** tab of the Properties dialog box for an adapter or staging transformation.

ASG TMONC2CIC Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file	any valid file	none	
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Accept Discard Terminate	Discard	

Parameters Supported for This Adapter	Values	Default Value	Notes
Logical record length	any positive integer less than or equal to 32,760	0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.	

ASG TMONDB2 Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file	any valid file		
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Accept Discard Terminate	Discard	
Logical record length	positive number less than or equal to 32,760	0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.	

BMC Mainview IMS Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

BMC Perf Mgr Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file	any valid file		
Allow duplicate ID variables	Yes No		
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Accept Discard Terminate	Discard	

CA TMS Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

DT Perf Sentry Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file or directory	any valid file or directory		
Delimiter in raw data	Comma Tab	Comma	
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Accept Discard Terminate	Discard	

DT Perf Sentry with MXG Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

HP Perf Agent Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file or directory	any valid file or directory		
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Accept Discard Terminate	Discard	

HP Reporter Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input library: Library and Libref	any valid library and libref		
Future data	Accept Discard Terminate	Discard	
Default duration	0 or greater	0	

IBM DCOLLECT Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

IBM EREP Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

IBM SMF Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File or directory (default)	Raw data input file or directory	any valid file or directory	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
JES	JES2 JES3			JES2
Presummari- zation duration	any integer greater than or equal to 1 and less than or equal to 2,147,483,647			0
Use intermediate staging view	any valid staged table name			

IBM TPF Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

IBM VMMON Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Choose raw data input type	File (default)	Raw data input file	any valid file	
	FTP	Host	any valid host	
		Port	any valid port	21
		User	any valid user	
		Password	any valid password	
		External filename	any valid file	
		Tape	Yes No	No
		RCMD	any valid site	SITE RDW
		Debug	Yes No	No
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard
Logical record length	any positive integer less than or equal to 32,760			0, to indicate that the logical record length is dependent on the type of input file and operating system that is used.

MS SCOM Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input library: Library and Libref	any valid library and libref		
Future data	Accept Discard Terminate	Discard	

SAP ERP Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input library: Library and Libref	any valid library and libref		
Future data	Accept Discard Terminate	Discard	
Presummari-zation duration	any integer greater than or equal to 1 and less than or equal to 2,147,483,647	3600	

SAR Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file or directory	any valid file or directory		
Duplicate checking	Discard Force Inactive Terminate	Discard	

Parameters Supported for This Adapter	Values	Default Value	Notes
Future data	Accept Discard Terminate	Discard	

SNMP Staging Parameters

Parameters Supported for This Adapter	Values	Suboptions	Values for Suboptions	Default Values
Raw data input directory	any valid directory			
rsh host command	any valid command			
Use snmpwalk to gather character data	Yes	snmpwalk executable	any valid executable	
		HostFile for snmpwalk	any valid file	
		Community value for snmpwalk	any valid community name	public
	No (default)			
Choose access command	HP NNM / Netview (default)	snmpColDump executable	any valid executable	
	RRDTool	rrdtool executable	any valid executable	
		Number of days to load	a number from 1 to 365	2
Duplicate checking	Discard Force Inactive Terminate			Discard
Future data	Accept Discard Terminate			Discard

User-Written Staging Parameters

The following table shows the parameters that are available from the **Staging** option of the **Staging Parameters** tab in the Properties dialog box of the staging transformation.

Staging Parameters Supported	Values	Default Value	Notes
Raw data input file or directory	any valid file or directory		
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Discard Accept Terminate	Discard	

The following table shows the parameters that are available from the **Duplicate Checking** option of the **Staging Parameters** tab in the Properties dialog box of the staging transformation.

Staging Parameters Supported	Values	Default Value	Notes
IDVAR variable for %RMDUPCHK macro	name of the SAS variable that identifies the system or machine that generated the input data	MACHINE	
TIMESTMP variable for %RMDUPCHK macro	name of the SAS variable that contains the datetime stamp that uniquely identifies the time of the event or interval that is being recorded	DATETIME	
ENDFILE variable for %RMDUPCHK macro	name of the SAS variable that is used as the END= keyword for the SAS INFILE statement that reads the raw data	_eof	

Staging Parameters Supported	Values	Default Value	Notes
INT value for %RMDUPCHK macro	maximum time gap (or interval) that is to be allowed between the timestamps on any two consecutive records from the same system or machine.	00:29	<p>If the interval between the timestamp values exceeds the value of this parameter, then an observation with the new time range is created in the control data set. This is referred to as a gap in the data.</p> <p>The value for this parameter must be provided in the format hh:mm, where hh represents hours and mm represents minutes. For example, to specify an interval of 14 minutes, use INT=0:14. To specify an interval of 1 hour and 29 minutes, use INT=1:29.</p>
RANGES value for %RMDUPCHK macro	number of ranges of data that are expected	20	<p>It also includes the data that has already been encountered and stored in the control data set. Ranges are stored in memory slots and need to be pre-allocated. If the allocated memory slots are not sufficient to store the ranges that are needed, the job must be rerun.</p> <p>A gap in the data that is larger than the time specified by the INT parameter indicates that a second range is needed. Gaps in the data should be an exceptional event, but it is best to overestimate the number of ranges that are needed in order to avoid rerunning the job.</p>

Staging Parameters Supported	Values	Default Value	Notes
SYSTEMS value for %RMDUPCHK macro	number of systems and machines from which you expect to process data	9	This value includes the number of systems and machines that are in the new incoming data. It also includes the data that has already been encountered and stored in the control data set.
KEEP value for %RMDUPCHK macro	number of weeks for which control data is to be kept	3	Because this value represents the number of Sundays between two dates, a value of 2 (the default) results in a maximum retention period of 20 days.

The following table shows these parameters that are available from the **User-Written** option of the **Staging Parameters** tab in the Properties dialog box of the staging transformation. For a detailed description of the staging parameters for the User-written Staging transformation, see [“Staging Parameters Tab” on page 234](#).

User-Written Parameters Supported	Values	Default Value	Notes
User-written code	directory path and file specification		
Generate a macro variable with input table name	Yes No	Yes	
Generate macro variables with target table names	Yes No	Yes	
Generate a macro with target table options	Yes No	Yes	
Generate Tabinfo and Varinfo tables	Yes No	No	
Generate a macro to create target table ATTRIB statements	Yes No	Yes	

User-Written Parameters Supported	Values	Default Value	Notes
Generate a macro to create computed columns	Yes No	Yes	
Generate a macro to create column lists	Yes No	Yes	
Generate a macro to assign columns to missing	Yes No	Yes	

VMware Data Acquisition Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input library: Library and Libref	any valid library and libref		
Age limit for data acquisition table (in days)	a positive integer	7	

VMware vCenter Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input library: Library and Libref	any valid library and libref		
Future data	Accept Discard Terminate	Discard	

Web Log Staging Parameters

Parameters Supported for This Adapter	Values	Default Value	Notes
Raw data input file or directory	any valid file or directory		
Delimiter in raw data	Space Comma Tab	Space	
Duplicate checking	Discard Force Inactive Terminate	Discard	
Future data	Accept Discard Terminate	Discard	
Machine	any valid machine name		
Site name	any valid site name		
Presummari-zation duration	any integer greater than or equal to 1 and less than or equal to 2,147,483,647	3600	

Appendix 4

Jobs That the Adapter Setup Wizard Creates for Domain Categories

The following table lists the domain categories that the Adapter Setup wizard enables you to select for each adapter. This table also shows whether the Adapter Setup wizard generates aggregation jobs, information map jobs, and reporting jobs for the domain category.

Table A4.1 Adapter Domain Categories and Associated Jobs

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
ASG TMON2CIC	None	No	No
ASG TMONDB2	None	No	No
BMC Mainview IMS	None	No	No

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
BMC Perf Mgr	CacheWindows includes Patrol Cache performance measures for Windows.	Yes	No
	CpuUnix includes global CPU performance measures for UNIX.	Yes	Yes
	CpuWindows includes global CPU performance measures for Windows.	Yes	Yes
	DiskUnix includes disk performance measures for UNIX.	Yes	Yes
	FilesystemUnix includes Filesystem performance measures for UNIX.	Yes	No
	KernelUnix includes system kernel performance measures for UNIX.	Yes	No
	LogicalDiskWindows includes logical disk performance measures for Windows.	Yes	Yes
	MemoryUnix includes memory and paging performance measures for UNIX.	Yes	Yes
	MemoryWindows includes memory and paging performance measures for Windows.	Yes	Yes
	NetworkUnix includes network performance measures for UNIX.	Yes	No
	NetworkWindows includes network performance measures for Windows.	Yes	No
	PagingFileWindows includes paging file performance measures for Windows.	Yes	No
	PhysicalDiskWindows includes physical disk performance measures for Windows.	Yes	No
	ProcessUnix includes process performance measures for UNIX.	Yes	No
	SwapUnix includes swap performance measures for UNIX.	Yes	No
	SystemWindows includes system performance measures for Windows.	Yes	No

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
CA TMS	None	No	No
DT Perf Sentry	Exchange InfoStore includes storage measures.	Yes	No
	Exchange InfoStoreMailbox includes mailbox storage measures.	Yes	No
	Exchange InfoStorePrivate includes private storage measures.	Yes	No
	Exchange InfoStorePublic includes public storage measures.	Yes	No
	Exchange MailServer includes mail server measures.	Yes	No
	Server Disk includes disk performance measures.	Yes	Yes
	Server Network includes network performance measures.	Yes	No
	Server PagingFile includes paging file performance measures.	Yes	No
	Server PhysicalFile includes physical disk performance measures.	Yes	No
	Server ProcessorCpu includes processor CPU performance measures.	Yes	No
	Server System includes system performance measures.	Yes	Yes
DT Perf Sentry with MXG	None	No	No

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
HP Perf Agent	Disk Unix includes disk performance measures for UNIX.	Yes	Yes
	Disk Windows includes disk performance measures for Windows.	Yes	Yes
	LogicalSystem Unix includes logical system performance measures for UNIX.	Yes	No
	Network Unix includes network performance measures for UNIX.	Yes	No
	Network Windows includes network performance measures for Windows.	Yes	No
	ProcessorCpu Unix includes processor CPU performance measures for UNIX.	Yes	No
	ProcessorCpu Windows includes processor CPU performance measures for Windows.	Yes	No
	System Unix includes system performance measures for UNIX.	Yes	Yes
	System Windows includes system performance measures for Windows.	Yes	Yes

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
HP Reporter	Disk Unix includes disk performance measures for UNIX.	Yes	Yes
	Disk Windows includes disk performance measures for Windows.	Yes	Yes
	LogicalSystem Unix includes logical system performance measures for UNIX.	Yes	No
	Network Unix includes network performance measures for UNIX.	Yes	No
	Network Windows includes network performance measures for Windows.	Yes	No
	ProcessorCpu Unix includes processor CPU performance measures for UNIX.	Yes	No
	ProcessorCpu Windows includes processor CPU performance measures for Windows.	Yes	No
	System Unix includes system performance measures for UNIX.	Yes	Yes
	System Windows includes system performance measures for Windows.	Yes	Yes
IBM DCOLLECT	None	No	No
IBM EREP	None	No	No

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
IBM SMF	CICS includes CICS hourly summarized transaction activity and response time performance.	Yes	No
	DB2 Batch includes DB2 Batch environment activity and performance.	Yes	Yes
	DB2 Cics includes DB2 CICS environment activity and performance.	Yes	No
	DB2 Ddf includes DB2 DDF environment activity and performance.	Yes	No
	DB2 Global includes DB2 global activity and performance.	Yes	Yes
	JOBS includes batch job activity and performance.	Yes	No
	OMVS includes OMVS activity and performance.	Yes	No
	RMF CecLpar includes RMF LPAR performance by CEC.	Yes	Yes
	RMF Device includes RMF device performance.	Yes	Yes
	RMF GoalModeService includes RMF workload period performance such as CPU, memory, I/O, service, and response time.	Yes	No
	RMF Lpar includes RMF LPAR performance by system.	Yes	Yes
	RMF PagingAndSwapping includes legacy RMF system performance including memory, paging, and swapping.	Yes	No
	RMF PagingFile includes RMF Paging file performance.	Yes	No

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
IBM SMF	RMF Processor includes RMF CPU processor performance.	Yes	No
	RMF System includes RMF system performance such as CPU, memory, I/O, queuing, paging, and swapping.	Yes	Yes
	RMF Workload includes RMF workload performance using normalized structure such as CPU, memory, I/O, service, and response time.	Yes	Yes
	RMF WorkloadPeriod includes RMF workload and workload period performance.	Yes	No
	TSO includes TSO activity and performance.	Yes	No
IBM TPF	None	No	No
IBM VMMON	None	No	No

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
MS SCOM	Exchange Disk includes Exchange server disk performance measures.	Yes	No
	Exchange InfoStore includes storage measures.	Yes	Yes
	Exchange InfoStoreMailbox includes mailbox storage measures.	Yes	Yes
	Exchange InfoStorePublic includes public storage measures.	Yes	No
	Exchange MailServer includes mail server measures.	Yes	No
	Exchange Network includes Exchange server network performance measures.	Yes	No
	Exchange Paging File includes Exchange server paging file performance measures.	Yes	No
	Exchange PhysicalDisk includes performance measures for the physical disk on the Exchange server.	Yes	No
	Server Disk includes disk performance measures.	Yes	Yes
	Server Network includes network performance measures.	Yes	No
	Server Paging File includes paging file performance measures.	Yes	No
	Server PhysicalDisk includes performance measures for the physical disk on the Windows server.	Yes	No
	Server ProcessorCpu includes processor CPU performance measures.	Yes	No
	Server System includes system performance measures.	Yes	Yes

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
SAP ERP	Batch includes SAP batch job activity and performance measures.	Yes	No
	Host includes SAP hourly application server activity and performance measures.	Yes	No
	System includes SAP hourly system activity and performance measures.	Yes	No
	SystemHostTask includes SAP hourly system, application server, and task type activity and performance measures.	Yes	No
	Task includes SAP hourly task type and elapsed time goal activity and performance measures.	Yes	No
	Transaction includes SAP hourly transaction activity and performance measures.	Yes	No
	Workload includes SAP workload activity and performance measures.	Yes	No
SAR	Disk includes disk performance measures.	Yes	Yes
	Network includes network performance measures.	Yes	No
	ProcessorCpu includes processor CPU performance measures.	Yes	No
	System includes system performance measures.	Yes	Yes
SNMP	ICMP includes SNMP MIB2 ICMP data.	Yes	Yes
	IF includes SNMP MIB2 interface data.	Yes	Yes
	IPDevice includes SNMP MIB2 IP device data.	Yes	Yes
	System includes SNMP MIB2 system data.	Yes	Yes
	TCP includes SNMP MIB2 TCP data.	Yes	Yes
	UDP includes SNMP MIB2 UDP data.	Yes	Yes

Adapter	Domain Categories and Descriptions in the Adapter Setup Wizard	Aggregation and Information Map Jobs	Reporting Jobs
VMware vCenter	Guest Disk includes measures for VMware guest disk activity.	Yes	Yes
	Guest Network includes measures for VMware guest network activity.	Yes	No
	Guest Storage includes measures for VMware guest storage activity.	Yes	No
	Guest System includes measures for VMware guest system activity.	Yes	Yes
	Host Disk includes measures for VMware Host disk activity.	Yes	No
	Host Network includes measures for VMware host network activity.	Yes	Yes
	Host Storage includes measures for VMware host storage activity.	Yes	No
	Host System includes measures for VMware host system activity.	Yes	Yes
Web Log	Browser Activity includes browser information about visiting clients by machine, domain, and host.	Yes	No
	Server Activity includes information about Web site activity and status codes by machine and site name.	Yes	No
	Visitor Organization Activity includes information about Web site visitors by domain, host, organization, country, and state.	Yes	No
	Web Page Activity includes information about Web page activity by machine, site name, and page name.	Yes	No

For more information about the performance areas and staged tables that each domain category includes for reporting and analysis, see [“Overview of the SAS IT Resource Management Data Model”](#) on page 459.

Appendix 5

Duplicate-Data Checking

Duplicate-Data Checking Overview	391
What Is Duplicate-Data Checking?	392
Control Data Sets for Duplicate-Data Checking	392
Content of Permanent Control Data Sets	393
Duplicate-Data-Checking Macros	394
Implementing Duplicate-Data-Checking Macros	394
About Implementing Duplicate-Data-Checking Macros	394
Steps to Implement Duplicate-Data-Checking Macros	395
Steps to Prevent Implementation of Duplicate-Data-Checking Macros	395
Instructions for Implementing Duplicate-Data Checking	396
Overview	396
Implementation for Supplied Adapters That Are MXG Based	396
Implementation for Supplied Adapters That Are Not MXG Based	399
Implementation for User-Written Staging Code That Is Based on MXG Code . . .	403
Implementation for User-written Staging Code That Is Not Based on MXG Code	410
Implementation for SAP ERP, MS SCOM, HP Reporter, and VMware Adapters	414
Examples	414
Example 1: Processing Data in a DATETIME Range for Which a DATETIME and IDVAR Was Already Processed	414
Example 2: Processing Data for Two IBM SMF Files	415
Example 3: Using Macro Variables to Subset Data for HP Reporter, MS SCOM, and VMware Adapters	415

Duplicate-Data Checking Overview

SAS IT Resource Management provides a set of macros that enable you to control whether duplicate data is processed into the IT data mart. In this context, duplicate data is defined as data whose datetime stamp is within a range of data that has already been processed into the IT data mart for that machine or system.

Each of the duplicate-data-checking macros performs a specific task. Together, these macros set up and manage duplicate-data checking. The main purpose of the macros is to check your data and to prevent duplicate data from being processed into the IT data mart. However, sometimes it is necessary to process data in a datetime range for which a machine's or system's data was already processed. For example, you might need to

process data into a table that you did not use earlier or a table that you accidentally deleted. You can specify that the data is to be accepted even though it appears to be duplicate data.

This appendix describes how to set up duplicate-data checking for your type of data. First, this appendix includes a continuation of the overview and a brief introduction to the components in duplicate-data checking. Then, this appendix includes a separate set of instructions for implementing duplicate-data checking for each type of data.

What Is Duplicate-Data Checking?

As raw data is being read, one of the macros that performs duplicate-data checking reviews the datetime information in each record and stores the information in a SAS data set called a temporary control data set. Later, by using intermediate control data sets, another macro merges the information in the temporary control data set into one or more SAS data sets that are called permanent control data sets.

When additional data is processed into the IT data mart, the timestamps of the incoming data are compared with the datetime information in the permanent control data sets in order to determine whether the new data has already been processed. If it has, the duplicate data is handled in the way that you specify.

A duplicate-data report is printed in the SAS log after the data is read. The report describes how many records were read for each machine or system and how many duplicates were found, if any.

Note: The first time that you use the macros, the permanent control data sets have not been built, so the macro %RMDUPCHK cannot check the input records. Your data is not checked or rejected for duplicates, but the permanent control data sets are created and the datetime information for this data is saved to them. Data is checked only on datetime, although SMF data is also checked for the system name. (For example, if you try to add a new record type, but you have already read other record types from that adapter for that time period, the records will not be kept.) The duplicate-data report contains only a limited amount of information about your data.

Control Data Sets for Duplicate-Data Checking

Control data sets (temporary, intermediate, and permanent) are the basis on which duplicate data is rejected. Therefore, it is important that you understand how the control data sets are created, used, updated, and stored.

- temporary control data sets

When input data is processed, the %RMDUPCHK macro creates a temporary control data set, WORK._DUPCNTL, which stores information about the raw data from one or more adapters. Specifically, for each machine or system that generated data, the temporary control data set stores the datetime ranges and record counts in the raw data. If raw data from more than one data source is processed in the same job, then information is appended to the temporary control data set with each execution of the %RMDUPCHK macro.

- intermediate control data sets

When all the input data has been read and stored in the temporary control data set, the %RMDUPUPD macro writes the data from the temporary control data set, WORK._DUPCNTL, to separate intermediate data sets, which are located in the staging library for that staging job.

Note: The first time that a staging table is added to a staging transformation, the library is created that will contain the control data information for all the staged tables that are created by that staging transformation. It is called *adapter-nameStagingnnnn*, where *nnnn* is a random number that ensures that the library name is unique within the IT data mart. For example, a library name might be “DT Perf Sentry Staging 8926”. (This library also contains other types of data.) If all the staged tables that are created by that staging transformation are subsequently deleted, and one or more new staged tables are added to the transformation, then a new library is created for the new staged tables. The first library that was created (for the original staged tables) is not used again. It is not automatically deleted, but you can do so.

If the data set exists, then the data set is used and the new data overwrites the old data. Otherwise, the data set is created and the new data is written to that data set.

- permanent control data sets

The data in the intermediate control data sets is then merged by %RMDUPUPD into the corresponding permanent control data sets, which are named *sourceCNTRL* in the staging library for that staging job. The permanent control data sets are stored and maintained in the staging library. One permanent control data set, named *sourceCNTRL*, can exist for each adapter. Each data set contains information about that adapter's machines or systems, datetime ranges, and record counts.

If a *sourceCNTRL* data set exists, the new data is merged into it. Otherwise, the data set is created and the new data is written to that data set.

Content of Permanent Control Data Sets

The content of the permanent control data sets is based on intervals and ranges in your data. The INT= parameter flags datetime gaps. When you specify the INT= parameter on %RMDUPCHK, you define the maximum gap allowable between records in the same range. If the gap between datetimes for two consecutive records exceeds this value, then a new range is created.

A range is deleted when the end-of-range datetime value is older than the number of weeks that are specified on the KEEP= parameter on %RMDUPCHK. However, if your data is continuous, then you have only one range. Your control information is never aged out because the end-of-range datetime value is constantly extended by new datetime information.

Here are the ways that ranges are used with data that is continuous and data that is not continuous:

- If your data is continuous and does not have any datetime gaps that exceed the value of the INT= parameter, then your data always updates the same range. In this case, the permanent control data set contains one range for each unique value of the variable that is specified by the IDVAR= parameter. The values of that variable are typically the machine or system names from which the raw data originated.
- If your data is not continuous, then the permanent control data set contains multiple ranges for each unique value of the variable that is specified by the IDVAR=

parameter. Each range is prefixed with a value of the variable that is specified by the IDVAR= parameter.

Note: For the HP Reporter and VMware adapters, one set of control data sets is created for each table by default. You can create just one set of control data sets for these adapters. To do so, specify a three-character identifier in the SOURCE= parameter of the %RMDUPCHK macro. SAS IT Resource Management then prefixes that identifier to the control data set.

Duplicate-Data-Checking Macros

How and when you should use the duplicate-data-checking macros depends on the type of data that you are using. The duplicate-data-checking macros and short descriptions of the tasks that they perform are in the following list. Before you begin to use these macros, review the following section entitled “Implementing Duplicate-Data-Checking Macros.”

- %RMDUPINT loads macro definitions that are used by the other duplicate-data-checking macros.
- %RMDUPDSN generates the name (WORK._DUPCNTL) of the temporary SAS data set that will contain datetime ranges for the data that is being processed into the active IT data mart.
- %RMDUPCHK checks for duplicate data by examining timestamps on data that is being read by the staging code. This macro also writes to the temporary control data set.
- %RMDUPUPD updates the permanent control data sets with information from a temporary control data set by way of the intermediate control data sets.

Implementing Duplicate-Data-Checking Macros

About Implementing Duplicate-Data-Checking Macros

When a job is generated and deployed from a SAS IT Resource Management client, the duplicate-data property becomes a SAS macro variable. In this appendix, any reference to the duplicate-data or DUPMODE property of a staging transformation can also be interpreted to refer to the duplicate-data-checking macro variable.

Note: If you edit the generated source code to modify the mode for duplicate-data checking, then your changes to the code will be lost if the job is regenerated or redeployed. Although you can save the functionality of the source code to the local file system, the original source code is preserved and the mode for duplicate-data checking is not updated in the SAS Metadata Repository entries.

To implement duplicate-data-checking macros, you must first set up processing without duplicate-data checking. Afterward, you can choose whether to implement duplicate-data checking by specifying parameters in the staging transformation.

Steps to Implement Duplicate-Data-Checking Macros

The steps and considerations that are required for implementing duplicate-data checking vary according to whether you are using a staging transformation for a supported adapter or if you are using user-written staging code.

If you are using a staging transformation for a supported adapter, perform the following steps:

1. Double-click to open the job that contains the staging transformation for the adapter that you want to work with. Right-click the staging transformation, and from the drop-down list, select **Properties**. On the Properties dialog box, select the **Staging Parameters** tab. This tab lists the parameters that apply to the staging transformation.
2. Specify the **Duplicate checking** option as **FORCE**, **DISCARD**, or **TERMINATE**. The staging transformation arranges for the %RMDUPINT, %RMDUPDSN, %RMDUPCHK, and %RMDUPUPD macro calls to be invoked accordingly in the adapter-specific staging code.

The following list describes the active parameters for duplicate-data checking. The corresponding macro call values are included in parentheses.

- **FORCE**

If the %RMDUPCHK macro specifies FORCE=YES, “duplicate” data is allowed into the IT data mart. The “duplicate” data should be data that appears to the macros to be duplicate, but is data that you want to process anyway. For example, perhaps the “duplicate” data was processed into the IT data mart before but was accidentally deleted. Or the “duplicate” data is being processed into a table that was not in use in the IT data mart when the data was processed earlier.

- **TERMINATE**

If the %RMDUPCHK macro specifies TERM=YES, processing terminates if duplicate data is encountered.

- **DISCARD**

If the %RMDUPCHK macro does not specify FORCE=YES and does not specify TERMINATE=YES, duplicate data is discarded.

If you are using user-written staging code, perform the following steps:

1. Place the calls to the %RMDUPINT, %RMDUPDSN, %RMDUPCHK, and %RMDUPUPD macros in the staging code.
2. Set the values of the FORCE= and TERM= macros in the call to %RMDUPCHK. This step is required for the code to process.
3. Click **OK** to close the Properties dialog box for the staging transformation.

Steps to Prevent Implementation of Duplicate-Data-Checking Macros

If you are using a staging transformation for a supported adapter and you do not want to implement duplicate-data checking, perform the following steps:

1. Access the **Staging Parameters** tab for the staging transformation.

2. Specify the **Duplicate checking** option as INACTIVE.

Instructions for Implementing Duplicate-Data Checking

Overview

The instructions for implementing duplicate-data checking vary depending on how you have implemented processing. The following sections in this appendix provide detailed implementation instructions for various processing scenarios. See the following topics that correspond with how you are implementing processing:

- If you are using supplied adapters that are MXG based, then see [“Implementation for Supplied Adapters That Are MXG Based” on page 396](#).
- If you are using supplied adapters that are not MXG based, then see [“Implementation for Supplied Adapters That Are Not MXG Based” on page 399](#).
- If you are using user-written staging code that is based on MXG code, then see [“Implementation for User-Written Staging Code That Is Based on MXG Code ” on page 403](#).
- If you are using user-written staging code that is not based on MXG code, then see [“Implementation for User-written Staging Code That Is Not Based on MXG Code ” on page 410](#).

Note: Duplicate-data-checking macros are designed to prevent the same data from being processed into the IT data mart twice. However, there might be times when you need to backload data, that is, to process some data that is in a datetime range for which the permanent control data sets have already recorded machine or system data. (For example, you might need to process data into one or more tables that you did not use earlier, or to process data into one or more tables that you accidentally purged or deleted.) Remember to restore the DUPMODE=FORCE or FORCE=YES setting to its original value after you finish that task.

Implementation for Supplied Adapters That Are MXG Based

Overview

Use the following instructions if you process data with a supplied adapter that is based on MXG code. Currently, these adapters include IBM DCOLLECT, IBM EREP, BMC Mainview IMS, DT Perf Sentry with MXG, IBM SMF, ASG TMON2CIC, ASG TMONDB2, CA TMS, IBM TPF, or IBM VMMON.

Here is an overview of the preparation for duplicate-data checking for this case:

1. On the **Staging Parameters** tab of the staging transformation, set the **Duplicate checking** option to FORCE, DISCARD, or TERMINATE.
2. Create a member or file that contains a call to the %RMDUPCHK macro.
3. Review that member or file and modify it, if necessary. The staging transformation automatically inserts calls to the %RMDUPINT, %RMDUPDSN, %RMDUPCHK, and %RMDUPUPD macros.

The following sections describe the preceding steps in detail.

Set the Duplicate Checking Parameter on the Staging Parameters Tab

On the **Staging Parameters** tab of the staging transformation, set the **Duplicate checking** option to **TERMINATE**, **DISCARD**, or **FORCE**. Each of these values indicates that you do want to use duplicate-data checking. The particular value that you specify indicates how you want duplicate data to be handled.

TERMINATE

stops processing if duplicate data is encountered.

DISCARD

continues processing while rejecting duplicate data if it is encountered.

FORCE

continues processing and accepts duplicate data if it is encountered.

Note: Use this setting only when loading data that is not duplicate, even though it might be in the same datetime range and for the same machine or system as data that was loaded previously.

Create the Member or File

The MXG.SRCLIB is the primary MXG source library. The MXGUSER.SRCLIB directory is initially empty. It stores members that the user copies from MXSG.SRCLIB for customization.

Note: On z/OS, the MXGUSER.SRCLIB and the MXG.SRCLIB directories are created when SAS IT Resource Management is installed. If you are working with MXG on Windows or UNIX, you must install MXG separately and take note of its installed locations. This document will continue to use the terms MXG.SRCLIB, for the primary MXG source library, and MXGUSER.SRCLIB for the folder where you store your customizations.

The default call to the %RMDUPCHK macro is contained in a different member or file on each operating system. The following list shows how to enable duplicate-data checking for the IBM SMF adapter.

- On UNIX, copy the **dup_smf.sas** file from the SAS IT Resource Management **misc** directory to the MXGUSER.SOURCLIB directory. Keep the same filename.
- On Windows, copy the **dup_smf.sas** file from the SAS IT Resource Management **sasmisc** directory to the MXGUSER.SRCLIB directory. Keep the same filename.
- On z/OS, copy the **CPDUPSMF** member from the SAS IT Resource Management **CPMISC** PDS to the MXGUSER.SRCLIB PDS. Keep the same member name.

The following table lists the names of the members to copy from the z/OS **CPMISC** PDS. Files with similar names are in the **misc** directory on UNIX and the **sasmisc** directory on Windows.

Find your adapter and the version of SAS (such as 9.1.3, or later) that you are using. That combination of adapter and version has one or more members. Copy the member or members from the appropriate location to the MXGUSER.SRCLIB PDS or directory. Keep the same member names.

Table A5.1 Each Adapter's CMDUPxxx Members and CPDUPxxx Members

Adapter	Members
IBM DCOLLECT	CMDUPDCO
IBM EREP	CMDUPERP, CMDUPER1, CMDUPER2
BMC Mainview IMS	CMDUPIMF
DT Perf Sentry with MXG	CMDUPNTS
IBM SMF	CPDUPSMF
ASG TMONDB2	CMDUPTMD
ASG TMON2CIC	CMDUPTM2
CA TMS	CMDUPTMS
IBM TPF	CMDUPTPF
IBM VMMON	CMDUPVMM

Modify the Member or File

The new member (on z/OS) or new file (on UNIX or Windows) contains a call to the %RMDUPCHK macro. The call is similar to this one for IBM SMF:

```
%RMDUPCHK (SOURCE=SMF,
            IDVAR=SYSTEM,
            Timestmp=SMFTIME,
            ENDFILE=ENDOF SMF,
            INT=00:29,
            KEEP=9
            );
```

The following list describes the parameters and values that are used in this example:

- **SOURCE=identifier**
a three-character identifier of the data source or adapter. This value must be set to SMF for IBM SMF data. For a list of the SOURCE= names for the supported adapters, see “%RMDUPCHK Required Arguments” on page 514.
- **IDVAR=variable-name**
identifies the SAS variable that is used by MXG to denote the origin of each IBM SMF record. This value must be set to SYSTEM for the IBM SMF adapter, because that is the name of the variable that is used by MXG.
- **Timestmp=SAS-timestamp-variable**
identifies the SAS variable that is used by MXG to record the datetime stamp of each IBM SMF record. The variable name must be SMFTIME for the IBM SMF adapter, because that is the name of the timestamp variable that is used by MXG.

- **ENDFILE=variable-name**
identifies the name of the SAS variable that is used by MXG in the INFILE statement to denote the end-of-file condition for the IBM SMF data. The variable name must be ENDOFSMF for the IBM SMF adapter.
- **INT=interval**
represents the maximum interval that is permitted between the datetime stamps on any two consecutive IBM SMF records from the same z/OS system. If the interval between the datetime stamp values exceeds the value of this parameter, then a new range is created. For IBM SMF data, the default value for the INT= parameter is 29 minutes.
- **KEEP=number-of-weeks**
represents the maximum number of weeks for which you want to retain control data. A range is aged out (removed) when the end-of-range datetime is older than permitted by the value of this parameter. For IBM SMF, the default value is nine weeks.

For a complete description of these and other parameters for this macro, see [“%RMDUPCHK” on page 513](#).

Review and modify your adapter's default call as follows:

- Review the values of the INT= and KEEP= parameters, and change those values to be appropriate for your site.
- Do not use the TERM= parameter or FORCE= parameter. The default settings are TERM=NO and FORCE=NO. (If the **Duplicate checking** value on the **Staging Parameters** tab of the staging transformation is **TERMINATE**, then termination occurs when duplicate data is encountered and the setting of TERM=YES on %RMDUPCHK is not required. Likewise, if the **Duplicate checking** value on the **Staging Parameters** tab of the staging transformation is **FORCE**, then duplicate data is processed without requiring the setting of FORCE=YES on %RMDUPCHK.)
- Do not change any of the other parameters or values in your call.

Implementation for Supplied Adapters That Are Not MXG Based

Overview

Use the following instructions if you process data with a supplied adapter that is not based on MXG code and whose data does not reside in a database or in a SAS library. (Currently, the adapters that are not MXG based include HP Perf Agent, HP Reporter, DT Perf Sentry, BMC Perf Mgr, MS SCOM, SAP ERP, SAR, SNMP, VMware vCenter, VMware Data Acquisition, and Web Log.)

Here is an overview of the preparation for duplicate-data checking for this case:

1. If your data does not reside on a database, on the **Staging Parameters** tab of the staging transformation, set the **Duplicate checking** option to **FORCE**, **DISCARD**, or **TERMINATE**.
2. For any data for the above adapters, create an entry that contains a call to the %RMDUPCHK macro.
3. For any data for the above adapters, review that member or file and modify it, if necessary. The staging transformation automatically inserts calls to the %RMDUPINT, %RMDUPDSN, %RMDUPCHK, and %RMDUPUPD macros.

The following sections describe the preceding steps in detail.

Set the Duplicate Checking Parameter on the Staging Parameters Tab

For only those adapters whose data does not reside on a database, on the **Staging Parameters** tab of the staging transformation, set the **Duplicate checking** option to **TERMINATE**, **DISCARD**, or **FORCE**. Each of these values indicates that you do want to use duplicate-data checking. The particular value that you specify indicates how you want duplicate data to be handled.

TERMINATE

stops processing if duplicate data is encountered.

DISCARD

continues processing while rejecting duplicate data if it is encountered.

FORCE

continues processing and accepts duplicate data if it is encountered.

Note: Use this setting only when loading data that is not duplicate, even though it might be in the same datetime range and for the same machine or system as data that was loaded previously.

Note: The SNMP adapter requires that duplicate checking be turned on. This setting is necessary because neither method of gathering raw data for SNMP (HPNNM and rrdtool) can ensure that only the most recent raw data is being saved. Therefore, invoking the duplicate-data-checking code of SAS IT Resource Management is the only way to determine what is new data and what is duplicate data. For best results, the value for this parameter should always be set to **Discard**.

Create the Entry

The default call to the %RMDUPCHK macro is contained in the SASHELP library in a catalog entry named **SASHELP.ITMS.<adapter-name>_RMDUPCHK.SOURCE**.

The staging code looks for the call in the following location:

Table A5.2 Source Locations for Each Adapter

ADAPTER	%RMDUPCHK Location
HP Perf Agent	<staging-library>.HPOVPA.RMDUPCHK.SOURCE
HP Reporter	<staging-library>.HPOVREP.RMDUPCHK.SOURCE
DT Perf Sentry	<staging-library>.NTSMF.RMDUPCHK.SOURCE
BMC Perf Mgr	<staging-library>.PATROL.RMDUPCHK.SOURCE
SAP ERP	<staging-library>.SAP.RMDUPCHK.SOURCE
SAR	<staging-library>.SAR.RMDUPCHK.SOURCE
MS SCOM	<staging-library>.SCOM.RMDUPCHK.SOURCE
SNMP	<staging-library>.SNMP.RMDUPCHK.SOURCE
VMware vCenter	<staging-library>.VMWARE.RMDUPCHK.SOURCE

ADAPTER	%RMDUPCHK Location
Web Log	<staging-library>.WEBLOG.RMDUPCHK.SOURCE

To copy the default entry in SASHELP to the appropriate location in the staging library for an adapter, perform the following steps:

1. Start an interactive SAS session on the system where the SAS Workspace Server executes.
2. Allocate the staging library. You can determine the LIBNAME statement from SAS Management Console or SAS Data Integration Studio. To do so, right-click the staging library and select **View LIBNAME**. You can use the Copy and Paste functions to copy the LIBNAME statement that is subsequently displayed and use it in your SAS session.
3. Submit the following code from the SAS Program Editor window:

```
PROC CATALOG CAT=SASHELP.ITMS;
    COPY OUT=staging-library.adapter-name;

    SELECT adapter-name_RMDUPCHK.SOURCE;
    RUN;
    QUIT;

PROC CATALOG CAT=staging-library.adapter-name;

    CHANGE adapter-name_RMDUPCHK.SOURCE=RMDUPCHK.SOURCE;

    RUN;
    QUIT;
```

Note: In the preceding code, replace *staging-library* with the libref submitted above. In addition, replace *adapter-name* with the adapter mnemonic listed in the %RMDUPCHK Location column of the [Table A5.2](#) (for example HPOVPA, HPOVREP, and so on).

Note: If the staging code is executed and there is no such entry in the staging library, then the staging code copies the default entry that is in SASHELP into the appropriate location in the staging library and displays a warning in the SAS log that notifies you of this action.

Modify the Entry

The new catalog entry contains a call to the %RMDUPCHK macro. To see the call, issue the following SAS command from a SAS command line (on z/OS) or from the command field on the SAS GUI (on Windows) or SAS Toolbar (on UNIX):

```
NOTE staging-library.adapter-name.RMDUPCHK.SOURCE
```

The command opens a Notepad window where the contents of the source entry are displayed. For example, this is the call to %RMDUPCHK for SAP ERP data:

```
%RMDUPCHK
    (SOURCE=SAP,
    IDVAR=SYSHOST,
    TIMESTMP=DATETIME,
```

```

ENDFILE=DUPEND,
INT=00:20,
KEEP=52
);

```

The following list describes the parameters and values that are used in this example:

- **SOURCE=identifier**
is a three-character identifier of the data source or adapter. This value must be set to SAP for the SAP ERP data. For a list of the SOURCE= names for the supported adapters, see “%RMDUPCHK Required Arguments” on page 514.
- **IDVAR=variable-name**
identifies the SAS variable that is used to denote the origin of each SAP ERP record. The variable name must be SYSHOST for the SAP ERP adapter and must not be changed.
- **TIMESTMP=SAS-timestamp-variable**
identifies the SAS variable that is used to record the datetime stamp of each SAP ERP record. The variable name must be DATETIME for the SAP ERP adapter and must not be changed.
- **ENDFILE=variable-name**
identifies the name of the SAS variable that is used in the INFILE statement to denote the end-of-file condition for the SAP ERP data. The variable name must be DUPEND for the SAP ERP adapter and must not be changed.
- **INT=interval**
represents the maximum interval that is permitted between the datetime stamps on any two consecutive SAP ERP records from the same machine and host combination. If the interval between the timestamp values exceeds the value of this parameter, then a new time range is created.
- **KEEP=number-of-weeks**
represents the maximum number of weeks for which you want to retain control data. A range is aged out (removed) when the end-of-range datetime is older than permitted by the value of this parameter. For SAP ERP data, the default value is 52 weeks.

For a complete description of these and other parameters for this macro, see “%RMDUPCHK” on page 513.

Review and modify your adapter's default call as follows:

- Review the values of the INT= and KEEP= parameters, and change those values to be appropriate for your site.
- Do not use the TERM= parameter or FORCE= parameter. The default settings are TERM=NO and FORCE=NO. (If the **Duplicate checking** value on the **Staging Parameters** tab of the staging transformation is **TERMINATE**, then termination occurs when duplicate data is encountered and the setting of TERM=YES on %RMDUPCHK is not required. Likewise, if the **Duplicate checking** value on the **Staging Parameters** tab of the staging transformation is **FORCE**, then duplicate data is processed without requiring the setting of FORCE=YES on %RMDUPCHK.)
- Do not change any of the other parameters or values in your call.

- When you are finished, press the END key or issue the END command from the command line to save the change and exit from the Notepad window. Then, you can exit the SAS session.

Implementation for User-Written Staging Code That Is Based on MXG Code

Overview

Use these instructions if you process data with user-written staging code that is based on MXG code.

Note: The following example demonstrates how to modify the MXG code that processes IBM DCOLLECT data. This example is for illustrative purposes only, and it uses IBM DCOLLECT data as a sample. In actual use, IBM DCOLLECT is a supported adapter, and SAS IT Resource Management modifies the MXG code as needed to process IBM DCOLLECT data. Thus, if you are using IBM DCOLLECT data in a real scenario, no MXG customization is necessary and you can use the instructions in the section [“Implementation for Supplied Adapters That Are MXG Based” on page 396](#).

Here is an overview of the preparation for duplicate-data checking in this case:

1. Create a member or file for a call to the %RMDUPCHK macro.
2. In that member or file, add the %RMDUPCHK call and modify the call as necessary.
3. Insert calls and include a member or file:
 - Insert a call to the %RMDUPINT, macro in the MXG code.
 - Insert a call to the %RMDUPDSN macro in the MXG code.
 - Include into the MXG code the member or file that contains the call to the %RMDUPCHK macro.
 - Add a call to the %RMDUPUPD macro at the end of your code.

The following sections describe the preceding steps in detail.

Create the Member or File

The following list details three hosts and the name of the member or file that is required in MXGUSER.SRCLIB for each. If the appropriate member or file does not already exist in MXGUSER.SRCLIB, create it and specify the name that corresponds with your host system.

z/OS Specifics

Create member `$$GENxxx`.

UNIX Specifics

Create file `gen_xxx.sas`.

Windows Specifics

Create file `gen_xxx.sas`.

Note: The value `xxx` in the preceding examples represents a three-character identifier for your data source or adapter. You can choose any identifier as long as it is not one of the identifiers on the list of values for the SOURCE= parameter in [“%RMDUPCHK” on page 513](#). For example, a three-character identifier for IBM DCOLLECT might be “DCO.”

Modify the Member or File

In the member or file created or used in the preceding step, add a call to %RMDUPCHK.

Here is an example for IBM DCOLLECT:

```
%RMDUPCHK (SOURCE=DCO, IDVAR=DCUSYSID, TIMESTMP=DCUTMSTP,
           ENDFILE=EOF, INT=25:00,
           KEEP=9,
           FORCE=NO, TERM=NO
           );
```

The following list describes the parameters and values that are used in this example:

- **SOURCE=identifier**
specifies a three-character identifier of the data source or adapter. In this example, “DCO” is the value that is used for IBM DCOLLECT data. For a list of the SOURCE= names for the supported adapters, see [“%RMDUPCHK Required Arguments” on page 514](#).
- **IDVAR=variable-name**
specifies the name of the SAS variable that is used by MXG to denote the origin of each record. This example uses the value DCUSYSID, because that is the variable name that is used in VMA CDCOL for IBM DCOLLECT data.
- **TIMESTMP=timestamp-variable-name**
specifies the name of the SAS variable that is used by MXG that contains the timestamp of each IBM DCOLLECT record. The variable name must be set to DCUTMSTP for IBM DCOLLECT data, because that is the name used in VMA CDCOL. Note that every IBM DCOLLECT record will have the same timestamp because it represents the time when the IDCAMS IBM DCOLLECT facility actually output the records. As shown in this example, you should change the value of the INT= parameter to accommodate this difference.
- **ENDFILE=variable-name**
specifies the SAS variable that represents the end-of-file condition for your source data. For IBM DCOLLECT data, this value must be set to EOF, because that is the name of the SAS variable that is used by MXG in the INFILE statement in VMA CDCOL in order to indicate the end-of-file condition.
- **INT=interval**
represents the maximum interval that is allowed between the datetime stamps on any two consecutive IBM DCOLLECT records from the same z/OS system. If the interval between the datetime stamp values exceeds the value of this parameter, then a new range is created. In this example, the specified value is 25 hours, because the timestamp for each IBM DCOLLECT record is the same. In effect, this enables the daily IBM DCOLLECT facility to run up to one hour late each day.
- **KEEP=number-of-weeks**
represents the maximum number of weeks for which you want to retain control data. A range is aged out (removed) when the end-of-range datetime is older than allowed by the value of this parameter. In this example, the specified number of weeks is nine.
- **FORCE=NO | YES**

FORCE=NO indicates that duplicate data should not be processed into the IT data mart.

- TERM=NO | YES

TERM=NO indicates that processing should not stop if duplicate data is encountered.

For a complete description of these and other parameters and appropriate values for each adapter, see “%RMDUPCHK” on page 513.

Review and modify your calls using the following guidelines and recommendations:

- For the SOURCE=*value*, specify the three-character identifier that you chose earlier.
- For the ENDFILE=*value*, refer to the MXG staging code for your adapter.
- For the IDVAR=*value*, refer to the MXG staging code for your adapter.
- For the TIMESTMP=*value*, refer to the MXG staging code for your adapter.
- Review the INT= and KEEP= values and change them to be appropriate for your site.
- Review the TERM= and FORCE= values and change them to be appropriate for your purpose. (Note that the combination TERM=YES, FORCE=YES is not meaningful.)

Insert Calls and Include Member or File

This part of the implementation process involves inserting the call to %RMDUPINT, inserting the call to %RMDUPDSN, and including the member or file that contains the call to %RMDUPCHK. These instructions depend on whether you are using the IMACxxxx or IMACKEEP member to specify your data source customizations.

If you are using the IMAC member to specify your data source customizations, perform the following steps:

1. Insert a call to %RMDUPINT.
 - a. In MXG.SRCLIB, in member or file IMACAAAA, there is a table with the names of the other IMAC members or files and a description of what is in each. Find the name of the IMAC member or file for your adapter. For example, the member IMACDCOL is for IBM DCOLLECT data.
 - b. If your adapter's IMAC member or file is not already in MXGUSER.SRCLIB, copy that member or file from MXG.SRCLIB to MXGUSER.SRCLIB. Keep the same name.
 - c. At the beginning of your adapter's IMAC member or file in MXGUSER.SRCLIB, insert a call to the %RMDUPINT macro. Use the following form: **%RMDUPINT ;**

Note: The %RMDUPINT macro does not have any parameters.

2. Insert a call to %RMDUPDSN.

In that same IMAC member or file, inside any one (but only one) of the _L macro definitions, insert a call to the %RMDUPDSN macro. Use the following form: **%RMDUPDSN (SOURCE=xxx)** where xxx is a three-character identifier for your data source or adapter. You can choose any identifier as long as it is not one of the identifiers that is on the list of values for the SOURCE= parameter in the %RMDUPCHK macro.

Note: The call has an unusual form: there are two leading percent signs and no trailing semicolon.

For example, in IMACDCOL, you might change the following code:

```
MACRO _LDCODSN DCOLDSET %
MACRO _KDCODSN          %
MACRO _LDCOCLU DCOLCLUS %
...
```

The changed version would look like this:

```
MACRO _LDCODSN %%rmdupdsn(source=DCO) DCOLDSET %
MACRO _KDCODSN          %
MACRO _LDCOCLU DCOLCLUS %
...
```

3. Include the member or file that contains the call to %RMDUPCHK.
 - a. In member or file IMACAAAA of MXG.SRCLIB, there is a list of the names of the members or files and a description of what is in each. Find the name of the member or file that decodes data for your adapter. For example, the member or file for IBM DCOLLECT data is named VMACDCOL.
 - b. In MXG.SRCLIB, scan through your adapter's VMAC member or file for a reference to an IHDR member or file. For example, for IBM DCOLLECT data in MXG version 14.02 and later, VMACDCOL refers to a member or file that is called IHDRDCOL.
 - c. If the VMAC member or file does refer to an IHDR member or file, include the member or file that contains the call to %RMDUPCHK in the IHDR member or file.

If the member or file with the IHDR name is not already in MXGUSER.SRCLIB, copy that member or file from MXG.SRCLIB to MXGUSER.SRCLIB. Keep the same name.

At the beginning of that member or file in MXGUSER.SRCLIB (typically, the member or file is empty), insert the following code:

z/OS Specifics

```
%INCLUDE '<codeFocus>fully-qualified-PDS</codeFocus>($$GENxxx)' ;
```

UNIX Specifics

```
%INCLUDE '<codeFocus>full-path-directory</codeFocus>/gen_xxx.sas' ;
```

Remember that directory names and filenames are case sensitive on UNIX.

Windows Specifics

```
%INCLUDE full-path-directory\gen_xxx.sas' ;
```

For *fully-qualified-PDS* or *full-path-directory*, substitute the location of MXGUSER.SRCLIB. The member or filename refers to the member or file that you created in the preceding steps. The *xxx* refers to the three-character identifier that you chose in the preceding steps.

Note: The %INCLUDEs in the preceding examples use absolute locations. If you prefer to use relative locations, use the following code.

z/OS Specifics

```
%INCLUDE SOURCLIB($$GENxxx) ;
```

UNIX Specifics

```
%INCLUDE SOURCLIB(gen_XXX);
```

Windows Specifics

```
%INCLUDE SOURCLIB(gen_XXX);
```

- d. If the VMAC member or file does not refer to an IHDR member or file, include the member or file that contains the call to %RMDUPCHK in every one of the adapter's EX members or files.
 - Identify all of the EX members or files for your adapter. To find the member names, scan the adapter's VMAC member or file in MXG.SRCLIB and look for %INCLUDE statements that refer to members of SOURCLIB that have the prefix EX.
 - For each EX member or file, if a member or file with the same name does not already exist in MXGUSER.SRCLIB, copy the EX member or file from MXG.SRCLIB into MXGUSER.SRCLIB. Keep the same name.
 - In each of these EX members or files in MXGUSER.SRCLIB, there is an OUTPUT statement. In front of the OUTPUT statement, on a new line, insert the following line of code:

z/OS Specifics

```
%INCLUDE
'fully-qualified-PDS($$GENxxx)';
```

UNIX Specifics

```
%INCLUDE
'full-path-directory/gen_XXX.sas';
```

Remember that directory names and filenames are case sensitive on UNIX.

Windows Specifics

```
%INCLUDE
'full-path-directory\gen_XXX.sas';
```

- For *fully-qualified-PDS* and *full-path-directory*, substitute the location of MXGUSER.SRCLIB. The member name or filename refers to the member or file that you created in the preceding steps. The *xxx* refers to the three-character identifier that you chose in the preceding steps.

The %INCLUDEs in the preceding examples use absolute locations. If you prefer to use relative locations, use the following code:

z/OS Specifics

```
%INCLUDE SOURCLIB ($$GENxxx);
```

UNIX Specifics

```
%INCLUDE SOURCLIB (gen_XXX);
```

Windows Specifics

```
%INCLUDE SOURCLIB (gen_XXX);
```

Note: Most data sources have several OUTPUT statements (one OUTPUT statement per EX member) and thus the potential for calling %RMDUPCHK several times. The duplicate-data-checking macros have verification routines to ensure that %RMDUPCHK executes only once and checks the data only once.

4. After the end of the MXG code, add a call to the %RMDUPUPD macro. Use the following form: **%RMDUPUPD;**

Note: The %RMDUPUPD macro does not have any parameters.

If you are using the IMACKEEP member to specify your data source customizations, perform the following steps:

1. Insert a call to the %RMDUPINT macro.
 - a. If the IMACKEEP member or file is not already in MXGUSER.SRCLIB, copy it from MXG.SRCLIB to MXGUSER.SRCLIB. Keep the same name.
 - b. In the IMACKEEP member or file in MXGUSER.SRCLIB, insert a call to the %RMDUPINT macro immediately after the %INCLUDE statement for IMACOLDV. Use the following form: **%RMDUPINT;**

Note that the %RMDUPINT macro does not have any parameters.

2. Insert a call to the %RMDUPDSN macro.
 - a. In the IMACKEEP member or file in MXGUSER.SRCLIB, there is a group of `_W` macro definitions for your adapter. Copy any one (but only one) of these `_W` macro definitions and insert it immediately after the call to the %RMDUPINT macro.
 - b. Inside that copy, insert a call to the %RMDUPDSN macro. Use the following form: **%%RMDUPDSN (SOURCE=xxx)** where *xxx* is the three-character identifier that you chose in the preceding steps. *Note* that the call has an unusual form: there are two leading percent signs and no trailing semicolon. For example, for IBM DCOLLECT, you might copy the following macro definition:

```
MACRO _WDCODSN DCOLDSET%
```

When you insert the call, the copy will look like this:

```
MACRO _WDCODSN %%RMDUPDSN(SOURCE=DCO) DCOLDSET %
```

3. Include the member or file that contains the call to the %RMDUPCHK macro.
 - a. In MXG.SRCLIB, in member or file IMACAAAA, there is a list of the names of the members or files and a description of what is in each. Find the name of the member or file that decodes data for your adapter. For example, the member or filename for IBM DCOLLECT data is VMACDCOL.
 - b. Scan through your adapter's VMAC member or file for a reference to an IHDR member or file. For example, for IBM DCOLLECT data in MXG version 14.02 and later, VMACDCOL refers to a member or file that is called IHDRDCOL.
 - c. If the VMAC member or file does refer to an IHDR member or file, include the member or file that contains the %RMDUPCHK call into the IHDR member or file.
 - If the member or file with the IHDR name is not already in MXGUSER.SRCLIB, copy that member or file from MXG.SRCLIB to MXGUSER.SRCLIB. Keep the same name.
 - At the beginning of that member or file in MXGUSER.SRCLIB (typically, the member or file is empty), insert the following line of code:

z/OS Specifics

```
%INCLUDE 'fully-qualified-PDS($$GENxxx)';
```

UNIX Specifics

```
%INCLUDE 'full-path-directory/gen_XXX.sas';
```

Remember that directory names and filenames are case sensitive on UNIX.

Windows Specifics

```
%INCLUDE 'full-path-directory\gen_XXX.sas';
```

- For *fully-qualified-PDS* or *full-path-directory*, substitute the location of MXGUSER.SRCLIB. The member name or filename refers to the member or file that you created in the preceding steps. The *xxx* refers to the three-character identifier that you chose in the preceding steps.

Note: The %INCLUDEs in the preceding examples use absolute locations. If you prefer to use relative locations, use the following code:

z/OS Specifics

```
%INCLUDE SOURCLIB ($$GENxxx);
```

UNIX Specifics

```
%INCLUDE SOURCLIB (gen_XXX);
```

Windows Specifics

```
%INCLUDE SOURCLIB (gen_XXX);
```

- d. If the VMAC member or file does not refer to an IHDR member or file, include the member or file that contains the %RMDUPCHK call into _E macro definitions in the IMACKEEP member or file.

In your adapter's VMAC member or file in MXGUSER.SRCLIB, there is a group of _E macro definitions. Copy all of these _E macro definitions from the VMAC member or file to the beginning of the IMACKEEP member or file.

Inside each of the _E macro definitions in the IMACKEEP member or file, include the member or file that contains the call to %RMDUPCHK. Use the following form:

z/OS Specifics

```
%%INCLUDE 'fully-qualified-PDS ($$GENxxx)';
```

UNIX Specifics

```
%%INCLUDE 'full-path-directory/gen_XXX.sas';
```

Remember that directory names and filenames are case sensitive on UNIX.

Windows Specifics

```
%%INCLUDE 'full-path-directory>\gen_XXX.sas';
```

To specify the values for *fully-qualified-PDS* and *full-path-directory*, substitute the location of MXGUSER.SRCLIB. The member name or filename refers to the member or file that you created in the preceding steps. The *xxx* refers to the three-character identifier that you chose in the preceding steps.

Note: The %INCLUDEs in the preceding examples use absolute locations. If you prefer to use relative locations, use the following code:

z/OS Specifics

```
%%INCLUDE SOURCLIB ($$GENxxx);
```

UNIX Specifics

```
%%INCLUDE SOURCLIB(gen_XXX);
```

Windows Specifics

```
%%INCLUDE SOURCLIB(gen_XXX);
```

For example, you might want to use relative locations instead of the absolute locations in the following code:

```
MACRO _EDCODSN
  %%INCLUDE SOURCLIB(EXDCODSN);
%
```

The revised code that uses relative locations would look like this:

```
MACRO _EDCODSN
  %%INCLUDE SOURCLIB($$GENDCO);
  %%INCLUDE SOURCLIB(EXDCODSN);
%
```

Note: Although most data sources have several `_E` macros and, therefore, several `%RMDUPCHK` calls, the duplicate-data-checking macros use verification routines to ensure that the macro executes only once and checks the data only once.

Implementation for User-written Staging Code That Is Not Based on MXG Code

Overview

Use these instructions if you process data with user-written staging code that is not based on MXG code. Here is an overview of the preparation for duplicate-data checking in this case:

1. Create a member or file for a call to the `%RMDUPCHK` macro.
2. In that member or file, add the `%RMDUPCHK` call and modify the call as necessary.
3. Insert a call to the `%RMDUPINT` macro in your staging code.
4. Insert a call to the `%RMDUPDSN` macro in your staging code.
5. Include in your staging code the member or file that contains the call to the `%RMDUPCHK` macro and add a call to the `%RMDUPUPD` macro at the end of your code.

The following sections describe, in detail, the steps that you perform.

Create the Member or File

MXG has a convenient storage place for the member or file.

If you have MXG, use the following list to determine how to specify the name of the member or file that is required in `MXGUSER.SRCLIB` for the host where you will process the data. If the appropriate member or file does not already exist in `MXGUSER.SRCLIB`, create it and use the name that corresponds with your host system.

z/OS Specifics

Create member `$$GENxxx`.

UNIX Specifics

Create file `gen_xxx.sas`.

Windows Specifics

Create file `gen_xxx.sas`.

If you do not have MXG, then use a storage location that you name. Use the following list to determine that action that you must perform:

z/OS Specifics

Create a PDS and, in it, create a member named `$$GENxxx`.

UNIX Specifics

Create a directory and, in it, create a file named `gen_xxx.sas`.

Windows Specifics

Create a directory and, in it, create a file named `gen_xxx.sas`.

The `xxx` noted in the preceding examples is a three-character identifier for your data source or adapter. You can choose any identifier as long as it is not one of the identifiers on the list of values for the `SOURCE=` parameter in “[%RMDUPCHK](#)” on page 513.

Modify the Member or File

In this member or file, if there is no call to the `%RMDUPCHK` macro, add a `%RMDUPCHK` call that is similar to the following code:

```
%RMDUPCHK (
    SOURCE=ABC,
    IDVAR=SYSTEMID,
    TIMESTMP=DTIME,
    ENDFILE=LAST,
    INT=00:10,
    KEEP=53)
;
```

Note: If you are running on z/OS, you will find some additional `%RMDUPCHK` examples in the SAS IT Resource Management PDS that is named CPMISC.

The following list describes the parameters and values that are used in this example:

- `SOURCE=identifier`
specifies a three-character identifier for your data source or adapter. This example uses the value ABC.
- `IDVAR=variable-name`
specifies the name of the SAS variable that identifies the system or machine from which the data originated. This example uses the value SYSTEMID.
- `TIMESTMP=timestamp-variable-name`
specifies the name of the SAS variable that contains the timestamp for each input record. In this example, the value is DTIME.
- `ENDFILE=variable-name`
identifies the SAS variable that represents the end-of-file condition for your INFILE statement. In this example, the value is LAST.
- `INT=interval`

represents the maximum interval that is allowed between the datetime stamps on any two consecutive input records from the same system or machine. If the interval between the datetime stamp values exceeds the value of this parameter, then a new range is created. In this example, a range can contain gaps no larger than 10 minutes.

- **KEEP=number-of-weeks**

represents the maximum number of weeks for which you want to retain control data. A range is aged out (removed) when the end-of-range datetime is older than permitted by the value of this parameter. In this example, the specified value is 53 weeks.

For a complete description of these and other parameters for this macro, see “%RMDUPCHK” on page 513.

Review and modify your calls using the following guidelines and recommendations:

- For the **SOURCE=** value, specify the three-character identifier that you chose earlier.
- For the **ENDFILE=** value, refer to your staging code.
- For the **IDVAR=** value, refer to your staging code.
- For the **TIMESTAMP=** value, refer to your staging code.
- Review the **INT=** and **KEEP=** values and change them to be appropriate for your site.
- Review the **TERM=** and **FORCE=** values and change them to be appropriate for your purpose.

Note: The combination **TERM=YES**, **FORCE=YES** is not meaningful.

Insert a Call to %RMDUPINT

In your staging code, insert a call to the %RMDUPINT macro in front of the **DATA** statement that initiates the **DATA** step that reads your raw data. Use the following form:

```
%RMDUPINT;
```

Note: The %RMDUPINT macro does not have any parameters.

Insert a Call to %RMDUPDSN

For the call, use the following form:

```
%RMDUPDSN (SOURCE=xxx)
```

In this form, **xxx** is the three-character source name that you chose in the preceding steps for this data source.

Note: The call has an unusual form: there is no trailing semicolon.

The call to the %RMDUPDSN macro must be inserted between the word **DATA** and the semicolon character (;) in the **DATA** statement that initiates the **DATA** step that reads the raw data and writes the staged data. For example, here is a **DATA** step that reads raw data and writes the staged data to **GENLIB.STAGED**:

```
DATA GENLIB.STAGED;
...
OUTPUT GENLIB.STAGED;
```

You would change that call as follows:

```
DATA GENLIB.STAGED %rmdupdsn (source=xxx) ;
```



```
...
OUTPUT GENLIB.STAGED;
```

Note: Any OUTPUT statement to the data set or view to which the data is being staged (in this case, GENLIB.STAGED) must explicitly specify the data set or view. Therefore, do not use the OUTPUT statement in this form:

```
OUTPUT ;
```

Instead, use the OUTPUT statement in this form:

```
OUTPUT libref.data-set-or-view-name;
```

Include the Member or File That Contains the Call to %RMDUPCHK

Place the %INCLUDE statement in your staging code after the Timestmp and IDVAR variables have been read in from your raw data and before any logic that would execute an OUTPUT statement. Within that range, place the %INCLUDE statement as early in your code as possible to avoid any unnecessary processing of data that is discovered to be duplicate data.

Point the %INCLUDE statement to the member or file that you created in the preceding steps. If you want to specify absolute locations, insert the following line in your staging code and use the following form:

z/OS Specifics

```
%INCLUDE 'fully-qualified-PDS($$GENxxx)';
```

UNIX Specifics

```
%INCLUDE 'full-path-directory/gen_XXX.sas';
```

Note: Remember that directory names and filenames are case sensitive on UNIX.

Windows Specifics

```
%INCLUDE 'full-path-directory\gen_XXX.sas';
```

For fully-qualified-PDS or full-path-directory, substitute the location of MXGUSER.SRCLIB. The xxx refers to the three-character identifier that you chose in the preceding steps.

If you want to specify relative locations, use the following form:

- If you stored the %RMDUPCHK call in MXGUSER.SRCLIB, use the following code:

z/OS Specifics

```
%INCLUDE SOURCLIB($$GENxxx);
```

UNIX Specifics

```
%INCLUDE SOURCLIB(gen_XXX);
```

Operating Environment Information

```
%INCLUDE SOURCLIB(gen_XXX);
```

- If you did not store the %RMDUPCHK call in MXGUSER.SRCLIB, create a fileref for your storage location by using the SAS FILENAME statement:

```
FILENAME myfref 'fully-qualified-PDS-name-or-full-path-directory-name';
```

In the statement, *myfref* is the name that you give to this fileref. After you create the fileref, use the following code:

z/OS Specifics

```
%INCLUDE myfref ($$GENxxx)
```

UNIX Specifics

```
%INCLUDE myfref (gen_xxx)
```

Windows Specifics

```
%INCLUDE myfref (gen_xxx)
```

Note: Do not add the .sas extension on UNIX or Windows. The extension will be appended automatically on UNIX and Windows.

Implementation for SAP ERP, MS SCOM, HP Reporter, and VMware Adapters

The staging code for the SAP ERP, MS SCOM, HP Reporter, and VMware adapters is designed to read in only those rows in the input database that were not previously staged. Even though the staged table is overwritten each time a staging job is executed, a record is kept of the data that was read so that only new data is read when the staging job is run.

To accomplish this, the staging code creates control data sets in the staging library for that adapter. For all adapters, two control data sets are created for each staged table that is processed. These data sets are named *xxxDAY* and *xxxCNTRL*. These control data sets determine the raw data observations that should be staged. If the control data sets do not exist when the staging job is executed, then all of the data in the input database is staged.

The staging code for the SAP ERP, MS SCOM, HP Reporter, and VMware adapters always uses the control data sets so that only new data is staged. Therefore, the duplicate-data checking option is active by default for these adapters. All subsequent staging jobs that specify that adapter's staging library location use that library's control data sets.

Note: You can turn off the duplicate-data checking option by modifying the generated SAS code. For an example that shows how to do so, see [“Example 3: Using Macro Variables to Subset Data for HP Reporter, MS SCOM, and VMware Adapters”](#) on page 415.

Examples

Example 1: Processing Data in a DATETIME Range for Which a DATETIME and IDVAR Was Already Processed

This example describes how to process data into a table that was not used earlier, or if the data is in a table that was accidentally deleted.

To process data in a datetime range for which a particular DATETIME and IDVAR variable was already processed, specify the **FORCE=YES** parameter on the

%RMDUPCHK macro. This parameter specifies that the data is to be accepted, even though it appears to be duplicate data.

Note: The IDVAR= parameter identifies the SAS variable that is used by MXG to denote the origin of each IBM SMF record. This value must be set to SYSTEM for the IBM SMF adapter, because that is the name of the variable that is used by MXG.

Example 2: Processing Data for Two IBM SMF Files

This example describes how to process data for two separate IBM SMF files. In this example, a site has two separate IBM SMF contiguous data sets for the same IDVAR variable. It attempts to stage them into a supplied SAS IT Resource Management table. The first IBM SMF input file covers a datetime stamp range from 9:00 a.m. to 10:00 a.m., the second from 10:00 am to 11:00 a.m.

In addition, the following conditions apply to this example:

- The DUPMODE=DISCARD duplicate data parameter is set for both jobs.
- The records that describe the last IBM SMF interval are split across the two SMF data sets.
- Both of these IBM SMF files are processed for the same IT data mart.
- A 15-minute IBM SMF interval is being used—for example, an interval that starts at 9:45 a.m. and ends at 9:59:59 a.m. The RMF 70-79 records that describe all the statistics in this same interval are partially written on the first data set. When the SYS1.MAN1 data set is full, the RMF 70-79 records that start at 9:59:59 a.m are written in the SYS1.MAN2 data set.
- Each time an IBM SMF file is read, the data is processed into SAS IT Resource Management and aggregated.

There are two methods that you can use to process this data:

- Load all IBM SMF data in a single job, concatenating the two IBM SMF data sets. This is the preferred method, because no data is lost.
- Load IBM SMF data in two unique JCL steps or batch jobs. Each job should read one IBM SMF file at a time.

Note: In each method, the first IBM SMF data set is loaded into the IT data mart. The duplicate-data-checking macros mark the data as correctly loaded from 9:45:00 a.m. to 9:59:59 a.m. However, because the second method processes the SMF data files individually, those records that match the DATETIME and IDVAR parameters that range from 9:45:00 a.m. to 9:59:59 a.m. are considered to be duplicates. Therefore, they are discarded.

Example 3: Using Macro Variables to Subset Data for HP Reporter, MS SCOM, and VMware Adapters

The staging transformations for the HP Reporter, MS SCOM, and VMware adapters do not provide a parameter for duplicate-data checking. Duplicate-data checking is always enabled for these adapters. SAS IT Resource Management provides two macro variables that enable you to subset data for the HP Reporter, MS SCOM, and VMware adapters: ITRM_LoadFromDate and ITRM_LoadToDate. These macro variables enable you to override the default action of subsetting the incoming data based on the duplicate-data control data sets.

The ITRM_LoadFromDate and ITRM_LoadToDate macro variables can be used in the following situations:

- to backload data into tables that are added to a staging job after it has already run once against a given set of data

Note: Only the data for HP Reporter and VMware adapters can be backloaded. MS SCOM data cannot be backloaded.

- to specify a datetime range that will be used during staging to extract only the data from the input database whose datetime stamps fall within the range specified

Note: When the ITRM_LoadFromDate and ITRM_LoadToDate macro variables are set, the duplicate data checking code is still executed. SAS IT Resource Management will discard any data that is detected as duplicates.

The following code sets the ITRM_LoadFromDate and ITRM_LoadToDate macro variables to valid start and end datetime values. These values will be used to subset the data from the database instead of the ranges in the duplicate data control data sets. This code should be added to the generated code or to the deployed job code for the staging job.

```
%let ITRM_loadFromDate=14FEB2010:00:00:00;  
%let ITRM_loadToDate=15FEB2010:23:59:00;
```

Appendix 6

Adapter Key Metrics

What Are Key Metrics?	417
Reporting Areas or Domain Categories	418
About Reporting Areas or Domain Categories	418
Server Performance Domain Category	419
Storage Usage Domain Category	426
Web Log Performance Domain Category	426
Workload Profiling and Characterization Domain Category	427
Network Performance Domain Category	429
SMF Adapter Key and Ranked Metrics for the System Domain Category	432
SMF Adapter Key and Ranked Metrics for the Lpar Domain Category	435
SMF Adapter Key and Ranked Metrics for the Goal Mode Service Domain Category	438
SMF Adapter Key and Ranked Metrics for the Workload Domain Category . . .	441
SMF Adapter Key and Ranked Metrics for the Workload Period Domain Category	443
SMF Adapter Key and Ranked Metrics for the Device Domain Category	446
SMF Adapter Key and Ranked Metrics for the Paging and Swapping Domain Category	447
SMF Adapter Key and Ranked Metrics for the CPU Processor Domain Category	448
SMF Adapter Key and Ranked Metrics for the Paging File Domain Category . . .	448
SMF Adapter Key and Ranked Metrics for the Batch Job, TSO and OMVS Session Performance Domain Category	449
SMF Adapter Key and Ranked Metrics for the CICS Transaction Activity Domain Category	452
SMF Adapter Key and Ranked Metrics for the DB2 Thread Accounting Domain Category	455

What Are Key Metrics?

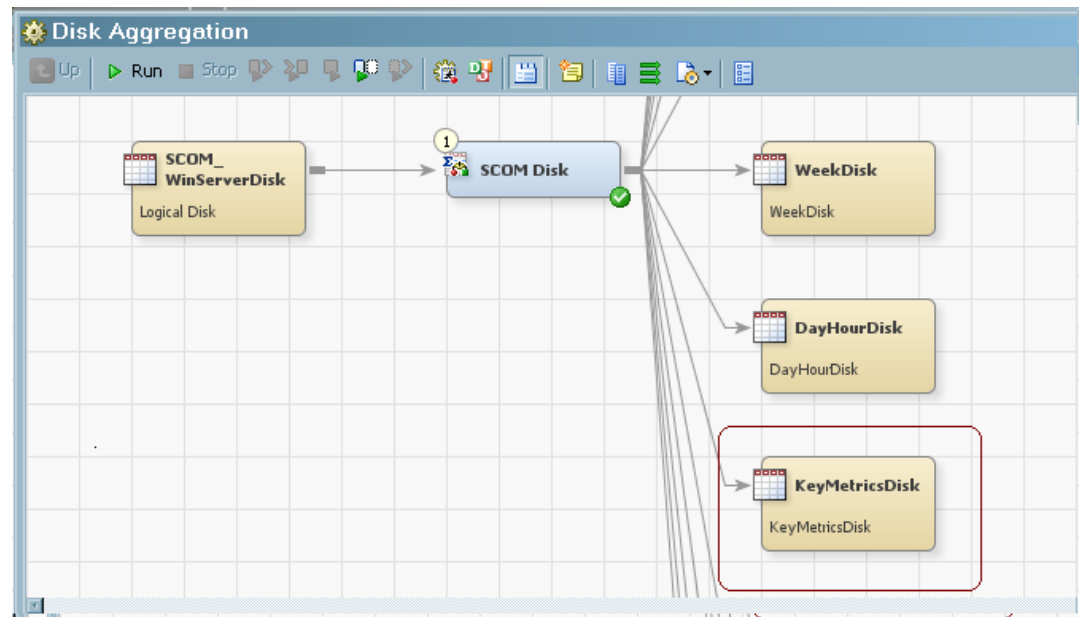
Key metrics are a subset of the available metrics provided by adapters for a given performance area such as resource utilization. SAS IT Resource Management identifies key metrics because monitoring hundreds of available metrics within and across adapters is not effective. You can use key metrics to quickly view the most critical data for a performance area. You can also collect key metrics over a long period of time to create capacity planning and forecasting plans.

Key metrics are standardized across adapters, enabling metrics from various data sources to be identified and compared appropriately. SAS IT Resource Management supports the

collection of these performance measurements in the IT data mart by staging and aggregating key metrics via adapter domain categories in the Adapter Setup wizard.

The Adapter Setup wizard creates key metric aggregations and regular aggregations for the specific domain categories that you choose for a given adapter. Key metrics aggregations are named with a prefix of “KeyMetrics.” The following image shows a key metrics aggregation, circled in red, in a process flow diagram.

Display A6.1 Key Metrics Aggregation in a Process Flow Diagram



Key metrics aggregations are designed to be narrow tables that serve as the source for capacity planning and forecasting. Key metrics aggregations include only key metrics and any relevant classification and ID columns that are associated with these key metrics to provide vital data for the performance metric.

Note: The measurements in a key metrics aggregation are considered key metrics and they are not ranked. However, these same individual metrics can be included in aggregations that are not key metrics aggregations. When an individual key metric is in a regular aggregation, it is ranked and is considered a ranked metric (not a key metric) in the context of the regular aggregation.

This appendix provides information about the key metrics that are provided by SAS IT Resource Management adapters. For more information about all the aggregation metrics that are provided by SAS IT Resource Management adapters, see the documentation at <http://support.sas.com/documentation/onlinedoc/itsv/metrics.html>.

Reporting Areas or Domain Categories

About Reporting Areas or Domain Categories

Adapters extract and transform data from various data sources to report on several performance metrics. These metrics can be categorized into the following logical reporting areas or domain categories:

server performance

metrics that are based on interval data collection in a server environment.

resource utilization

metrics showing the event-based activity for the resources associated with transactions, jobs, tasks, sessions, processes, and other units of work.

workload profiling and characterization

metrics based on interval or event-based data according to the methodology for measuring workload.

network performance and utilization

metrics including counts of network transmissions, device and interface utilization, device and interface status changes, and connection state changes.

metrics for the SMF adapter

key and ranked metrics including measurements for systems, Lpar, goal mode service, workload, workload period, device, paging and swapping, CPU processor, paging file, batch job, TSO, and OMVS sessions performance, CICS transaction activity, and DB2 thread accounting.

Server Performance Domain Category

Measuring server performance includes performance areas such as CPU, memory, I/O, paging and swapping, cache, and network time.

The following table shows the key metrics that are relevant to server performance and the adapter domain categories that support each metric. An X indicates that the key metric listed in the row is supported by the adapter domain category in the corresponding column.

Note: For the SAR, BMC Perf Mgr UNIX, and HP Perf Agent and HP Reporter UNIX domain categories, an X indicates that the corresponding metric is present for at least one of the supported environments (HP, Sun, AIX, Linux) that can run on the platform. The X does not necessarily mean that the metric is available in all environments for the domain category. For example, CpuWaitIoBusyPct is found in AIX and Sun environments on UNIX, but it is not available in HP and Linux environments.

In the following table, the abbreviated column headings stand for the following adapters:

- DT: Demand Technology Performance Sentry
- RMF: IBM Resource Management Facility
- SAR: System Activity Reporter
- BMCU: BMC Performance Manager for Servers - UNIX
- BMCW: BMC Performance Manager for Servers - Windows
- HPU: HP Perf Agent Server and HP Reporter Server for UNIX systems.
- HPW: HP Perf Agent Server and HP Reporter Server for Windows systems.
- SC: Microsoft System Center Operations Manager
- VMC: VMware vCenter Server

Table A6.1 Key Metrics for Server Performance Adapters

Metric Category	Column Label Root	DT	RMF	SAR	BMCU	BMCW	HPU	HPW	SC	VMC
CPU	AllEngCpuHardwareMips Usage		X							
CPU	AllEngCpuHardwareMsu Usage		X							
CPU	CpuBusyPct	X	X	X	X	X	X	X		X
CPU	CpuContextSwitchRate								X	
CPU	CpuIoInterruptRate	X	X	X	X	X	X	X		
CPU	CpuLoadAvg			X	X		X			
CPU	CpuReadyPct									X
CPU	CpuSystemBusyPct	X		X	X	X	X	X		
CPU	CpuUserBusyPct	X		X	X	X	X	X		
CPU	CpuWaitIoBusyPct			X	X		X			
CPU	FourHourMsuRollingAvg		X							
CPU	IfaEngCpuBusyPct		X							
CPU	IfaEngCpuHardwareMips Usage		X							
CPU	IfaEngCpuHardwareMsu Usage		X							
CPU	IflEngCpuBusyPct		X							
CPU	LogicalSystemCpuBusyPct						X			
CPU	LogicalSystemEntitledCpu						X			
CPU	LogicalSystemEntitledCpuBusyPct						X			
CPU	LogicalSystemHypervisor BusyPct						X			
CPU	LogicalSystemPhysicalCpuBusyPct						X			

Metric Category	Column Label Root	DT	RMF	SAR	BMCU	BMCW	HPU	HPW	SC	VMC
CPU	LogicalSystemPhysicalCpusUsage						X	X		
CPU	LogicalSystemRunQueueLength						X			
CPU	LparCpuBusyPct		X							
CPU	LparMsuHourlyRate		X							
CPU	MsuHourlyRate		X							
CPU	ReadyTasksWaitingPct		X							
CPU	RunQueueLength	X		X	X	X	X	X	X	
CPU	StandardEngCpuBusyPct		X							
CPU	StandardEngCpuHardwareMipsUsage		X							
CPU	StandardEngCpuHardwareMsuUsage		X							
CPU	StandardEngCpuSoftwareMsuUsage		X							
CPU	ZipEngCpuBusyPct		X							
CPU	ZipEngCpuHardwareMipsUsage		X							
CPU	ZipEngCpuHardwareMsuUsage		X							
Device	DeviceActivePct		X							
Device	DeviceBusyPct		X							
Device	DeviceIoRate		X							
Device	DeviceQueueAvgLength		X							
Device	DeviceResponseAvgMSec		X							
Disk	DiskBusyPct	X		X	X	X	X	X		
Disk	DiskPhysicalIoRateKB						X	X		
Disk	DiskPhysicalReadRate						X	X		
Disk	DiskPhysicalReadRateKB						X	X		

Metric Category	Column Label Root	DT	RMF	SAR	BMCU	BMCW	HPU	HPW	SC	VMC
Memory	CacheReadHitPct	X								
Memory	LogicalSystemPhysicalMemoryUsageMB						X	X		
Memory	LogicalSystemMemoryUsagePct						X			
Memory	MemoryActiveGB									X
Memory	MemoryAvailable		X							
Memory	MemoryAvailableMB	X		X	X	X	X	X	X	
Memory	MemoryAvailablePct						X	X		
Memory	MemoryBallooningGB									X
Memory	MemoryCacheUsagePct						X			
Memory	MemoryCommittedPct								X	
Memory	MemoryConsumedGB									X
Memory	MemoryGrantedGB									X
Memory	MemoryOverheadGB									X
Memory	MemoryPagesAvailableAix			X						
Memory	MemoryPagesAvailableSunSco			X						
Memory	MemorySharedGB									X
Memory	MemorySwapInGB									X
Memory	MemorySwapOutGB									X
Memory	MemorySystemCacheUsagePct						X	X		
Memory	MemorySystemUsagePct						X	X		
Memory	MemoryTotalUsage		X							
Memory	MemoryUsagePct	X	X	X			X	X		X
Memory	MemoryUserUsagePct						X	X		
Memory	MemoryVirtualUsagePct						X			

Metric Category	Column Label Root	DT	RMF	SAR	BMCU	BMCW	HPU	HPW	SC	VMC
Memory	MemoryZeroGB									X
Memory	PagingBlockRate			X						
Memory	PagingFileUsagePct	X							X	
Memory	PagingInputRate	X				X	X	X	X	
Memory	PagingOutputRate	X				X	X	X	X	
Memory	PagingRate	X	X	X	X	X	X	X	X	
Memory	PagingReadRate								X	
Memory	PagingWriteRate								X	
Memory	ReadyTasksWaitingPct		X							
Memory	RemainingFreeMemoryGB									X
Memory	SwapRate		X							
Memory	SwapSpaceUsagePct			X	X		X	X		
Memory	TotalMemoryLiabilityGB									X
Memory	TpsRecoveryPct									X
Network	InterfaceUsagePct	X				X	X	X		
Network	NetworkPacketsReceivedCount									X
Network	NetworkPacketsSentCount									X
Network	NetworkPacketsTotalCount									X
Network	NetworkReceivedRateKB								X	X
Network	NetworkSentRateKB								X	X
Network	PacketsTotalRate	X		X	X	X	X	X		
Network	ReceivedRateKB	X							X	
Network	SentRateKB	X							X	
Network	ThroughputRateKB			X		X	X	X		

Metric Category	Column Label Root	DT	RMF	SAR	BMCU	BMCW	HPU	HPW	SC	VMC
Network	TotalRateKB								X	
System Logical Memory	LogicalSystemPhysicalMemoryUsagePct						X			

The following table shows the key metrics that are relevant to exchange performance and the adapter domain category that supports each metric. An X indicates that the key metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.2 Key Metrics for Exchange Performance Domain Category

Metric Category	Column Label Root	DT Perf Sentry Exchange	MS SCOM
Exchange	ClientLogons		X
Exchange	ClientLogonsActive		X
Exchange	DiskQueueAvgLength		X
Exchange	DiskSpaceAvailableGB		X
Exchange	DiskSpaceAvailablePct		X
Exchange	MemoryAvailableMB		X
Exchange	MessagesDeliveredLocalMailRate	X	X
Exchange	MessagesSentPerMinute		X
Exchange	MessagesSubmittedPerMinute		X
Exchange	MessagesTotalRate	X	X
Exchange	PagingFileUsagePct		X
Exchange	PagingRate		X
Exchange	QueueLengthLocal	X	X
Exchange	QueueLengthRemote	X	X
Exchange	ReceivedRateKB		X
Exchange	ReceiveQueueSize	X	X
Exchange	RpcAvgLatencyMSec	X	X

Metric Category	Column Label Root	DT Perf Sentry Exchange	MS SCOM
Exchange	RpcClientTotalRateKB	X	
Exchange	RpcOperationsRate		X
Exchange	RpcPacketsRate	X	X
Exchange	RpcRequests		X
Exchange	RpcsSucceeded		X
Exchange	RunQueueLength		X
Exchange	SendQueueSize		X
Exchange	SentRateKB		X
Exchange	ThroughputRateKB	X	
Exchange	VmLargestBlockMB		X
Exchange	VmTotalLargeAvailableBlockMB	X	X

Storage Usage Domain Category

Measuring storage usage includes measuring capacity, free space, and used space of a virtual system.

The following table shows the key metrics that are relevant to storage usage and the VMware adapter domain categories that support this data. An X indicates that the key metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.3 Key Metrics for Storage Usage for VMware

Metric Category	Column Label Root	VMware
Storage	Capacity	X
Storage	FreeSpace	X
Storage	UsedSpace	X

Web Log Performance Domain Category

The following table shows the key metrics that are relevant to storage usage and the Web Log adapter domain categories that support this data. An X indicates that the key metric

listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.4 Key Metrics for Storage Usage for Web Log

Metric Category	Column Label Root	Web Log
Server	Pages	X
Server	Requests	X
Visitor	Pages	X
Visitor	Requests	X
WebPage	Pages	X
WebPage	Requests	X

Workload Profiling and Characterization Domain Category

Measuring workload profiling includes measuring the usage of CPU, memory and I/O resources, and the calculation of response time for individual workloads. Workloads are characterized by distinct classifications that are determined by the user.

The following table shows the key metrics that are relevant to workload and the adapter domain categories that support this data. An X in the column indicates that the key metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.5 Key Metrics for Workload Profiling and Characterization

Metric Category	Column Label Root	RMF Goal Mode Service	RMF Workload	RMF Workload Period	SAP ERP Workload
Workload/Service	AllEnginesCpuMipsUsage	X	X	X	
Workload/Service	CpuBusyPct		X	X	
Workload/Service	CpuCaptureTime	X			X
Workload/Service	CpuMipsUsage	X	X	X	
Workload/Service	CpuSrbServiceMsu	X			
Workload/Service	CpuTcbServiceMsu	X			
Workload/Service	DatabasePhysicalAccesses				X
Workload/Service	DatabasePhysicalChanges				X

Metric Category	Column Label Root	RMF Goal Mode Service	RMF Workload	RMF Workload Period	SAP ERP Workload
Workload/Service	DatabaseRecordRequests				X
Workload/Service	DatabaseRequests				X
Workload/Service	DatabaseRequestTime				X
Workload/Service	ElapsedTime				X
Workload/Service	ExcpRate		X	X	
Workload/Service	IfaCpuBusyPct		X	X	
Workload/Service	IfaCpuMipsUsage	X	X	X	
Workload/Service	IoRate		X	X	
Workload/Service	IoServiceMsu	X			
Workload/Service	MemoryUsage		X	X	
Workload/Service	MemoryServiceMsu	X			
Workload/Service	PerformanceIndex	X			
Workload/Service	ResponseTime		X	X	X
Workload/Service	ServiceMsu	X	X	X	
Workload/Service	ServiceRate	X	X	X	
Workload/Service	Total Storage Usage	X			
Workload/Service	Transactions		X	X	
Workload/Service	TransAvgResponse	X			
Workload/Service	TransCompleted	X			
Workload/Service	TransCompletedRate	X			
Workload/Service	TransCount				X
Workload/Service	TransferredKB				X
Workload/Service	TransMetGoalPct	X			
Workload/Service	TransMissedGoalPct	X			
Workload/Service	WaitTime				X

Metric Category	Column Label Root	RMF Goal Mode Service	RMF Workload	RMF Workload Period	SAP ERP Workload
Workload/Service	ZipCpuBusyPct		X	X	
Workload/Service	ZipCpuMipsUsage	X	X	X	

Note: In the preceding table, SAP ERP Workload includes all domain categories for the SAPR3 adapter.

Network Performance Domain Category

Measuring network performance includes measuring counts of network transmissions, device and interface utilization, device and interface status changes, and connection state changes.

The following table shows the key metrics that are relevant to network performance and the adapter domain categories that support this data. An X indicates that the key metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.6 Key Metrics for Network Performance

Metric Category	Column Label Root	SNMP
Network	IfIncomingOctet	X
Network	IfIncomingBit	X
Network	IfIncomingUnicastPkt	X
Network	IfIncomingNonUnicastPkt	X
Network	IfIncomingMulticastPkt	X
Network	IfIncomingBroadcastPkt	X
Network	IfIncomingPkt	X
Network	IfIncomingDiscardPkt	X
Network	IfIncomingErrorPkt	X
Network	IfIncomingUnkProtoPkt	X
Network	IfIncomingSuccessPkt	X
Network	IfInBitsPerPkt	X
Network	IfOutgoingOctet	X

Metric Category	Column Label Root	SNMP
Network	IfOutgoingBit	X
Network	IfOutgoingUnicastPkt	X
Network	IfOutgoingNonUnicastPkt	X
Network	IfOutgoingMulticastPkt	X
Network	IfOutgoingBroadcastPkt	X
Network	IfOutgoingPkt	X
Network	IfOutgoingDiscardPkt	X
Network	IfOutgoingErrorPkt	X
Network	IfOutgoingUnkProtoPkt	X
Network	IfOutgoingSuccessPkt	X
Network	IfUnicastPkt	X
Network	IfNonUnicastPkt	X
Network	IfMulticastPkt	X
Network	IfBroadcastPkt	X
Network	IfIncomingUtilization	X
Network	IfOutgoingUtilization	X
Network	IfUtilization	X
Network	IfAvailable	X
Network	IfStatusChange	X
Network	IpIncomingReceive	X
Network	IpIncomingHeaderError	X
Network	IpIncomingNoRoute	X
Network	IpIncomingAddressError	X
Network	IpIncomingUnkProto	X
Network	IpIncomingTruncDgram	X
Network	IpIncomingDiscard	X

Metric Category	Column Label Root	SNMP
Network	IpIncomingDeliver	X
Network	IpOutgoingRequest	X
Network	IpOutgoingNoRoute	X
Network	IpOutgoingTransmit	X
Network	IpOutgoingTransmit	X
Network	DeviceUp	X
Network	TcpAttemptFailure	X
Network	TcpEstablishedReset	X
Network	TcpCurrentConnEstablished	X
Network	TcpIncomingSegment	X
Network	TcpOutgoingSegment	X
Network	TcpRetransmitSegment	X
Network	TcpIncomingError	X
Network	TcpOutgoingRst	X
Network	TcpSegment	X
Network	IcmpIncomingMsg	X
Network	IcmpIncomingError	X
Network	IcmpOutgoingMsg	X
Network	IcmpOutgoingError	X
Network	IcmpError	X
Network	UdpIncomingDgram	X
Network	UdpNoPort	X
Network	UdpIncomingError	X
Network	UdpOutgoingDgram	X

SMF Adapter Key and Ranked Metrics for the System Domain Category

The following table shows the ITRM tables that contain the key and ranked metrics that are relevant to system domain categories and the SMF adapter that support this data.

Table A6.7 Tables for the SMF System Domain Category

ITRM Table	ITRM Table Description	MXG Dataset Name	ASW Available Table	Comments
XRMFSYS	RMF System	RMFINTRV TYPE71	Yes	Super table composed of the merging of subsets of metrics from MXG data sets RMFINTRV and TYPE71 to contain a snapshot view of important RMF system wide metrics.
XRMFINT	RMF Interval	RMFINTRV	No	Older legacy table present in ITRM releases before ITRM 3.2

The following table shows the key and ranked metrics that are relevant to system domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.8 Columns for the System Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PCMVSBY	MvsCpuBusyPct	WMean	X	X	
		WPStd	X		Used in forecasting **
PCCPUBY	CpuBusyPct (Standard Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
PCIFABY	IfaCpuBusyPct (zAAP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
PCZIPBY	ZipCpuBusyPct (zIIP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
TOTMEMR	MemoryTotalUsage	Mean	X		
		PStd	X		Used in forecasting **

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
MEMUSAGEPCT	MemoryUsagePct	WMean	X	X	
		WPStd	X		Used in forecasting **
MEMAVAIL	MemoryAvailable	Mean	X		
		Min	X	X	
		PStd	X		Used in forecasting **
IORATE7	IoInterruptRate	WMean	X	X	
		WPStd	X		Used in forecasting **
SYSMSU	SystemCpuMsuUsage (Standard Engine Hardware Msu)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the SYSMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
SYSMIPS	SystemCpuMipsUsag e (Standard Engine Hardware Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the SYSMSU column because these columns are ranked in a corresponding manner
IFAMSU	IfaCpuMsuUsage (zAAP Engine Hardware Msu)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the IFAMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
IFAMIPS	IfaCpuMipsUsage (zAAP Engine Hardware Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the IFAMSU column because these columns are ranked in a corresponding manner
ZIPMSU	ZipCpuMsuUsage (zIIP Engine Hardware Msu)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the ZIPMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
ZIPMIPS	ZipCpuMipsUsage (zIIP Engine Hardware Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the ZIPMSU column because these columns are ranked in a corresponding manner
TOTMSU	AllEngCpuMsuUsage (Hardware Msu for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the TOTMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
TOTMIPS	AllEngCpuMipsUsage (Hardware Mips for all Engines)	WMean	X		

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the TOTMSU column because these columns are ranked in a corresponding manner
MSUINTR	SystemMsuUsage (Standard Engine Software Msu)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	
PCRDYWT	ReadyTasksWaitingPct	WMean	X	X	
		WPStd	X		Used in forecasting **
PAGRT	PagingRate	WMean	X	X	
		WPStd	X		Used in forecasting **
SWAP	SwapRate	WMean	X	X	
		WPStd	X		Used in forecasting **

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the Lpar Domain Category

The following table shows the ITRM tables that contain the key and ranked metrics that are relevant to Lpar domain categories and the SMF adapter that support this data.

Table A6.9 Tables for the Lpar Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XASM70L	RMF Lpar	XASM70LP	Yes	LPAR measurements by Sysplex and System
XASMCEL	RMF CEC Lpar	XASMCELP	Yes	LPAR measurements by CEC

The following table shows the key and ranked metrics that are relevant to Lpar domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.10 Columns for the Lpar Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PCTLPBY	LparCpuBusyPct (Standard Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
PCIFABY	LparIfaCpuBusyPct (zAAP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
PCZIPBY	LparZipCpuBusyPct (zIIP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
PCIFLBY	LparIflCpuBusyPct (IFL Engine)	WMean	X		
		WPStd	X		Used in forecasting **
LPRSYMSU	LparCpMsuUsage (Standard Engine Hardware Msu)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the LPRSYSMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
LPRSYSMIPS	LparCpMipsUsage (Standard Engine Hardware Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the LPRSYMSU column because these columns are ranked in a corresponding manner
LPRIFAMSU	LparIfaMsuUsage (zAAP Engine Hardware Msu)	WMean	X		

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the LPRIFAMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
LPRIFAMIPS	LparIfaMipsUsage (zAAP Engine Hardware Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the LPRIFAMSU column because these columns are ranked in a corresponding manner
LPRZIPMSU	LparZipMsuUsage (zIIP Engine Hardware Msu)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		Although not ranked, the LPRZIPMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
LPRZIPMIPS	LparZipMipsUsage (zIIP Engine Hardware Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the LPRZIPMSU column because these columns are ranked in a corresponding manner
LPRTOTMSU	LparAllEngMsuUsage (Hardware Msu for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
		Sum	X		Although not ranked, the LPRTOTMIPS rankings can be used with this metric because these columns are ranked in a corresponding manner
LPRTOTMIPS	LparAllEngMipsUsage (Hardware Mips for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the LPRTOTMSU column because these columns are ranked in a corresponding manner
LPMSUHR	LparMsuHourlyRate	WMean	X	X	
		WPStd	X		Used in forecasting **
SM70LAC	LparFourHourAvg	WMean	X		
		WPStd	X		Used in forecasting **

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the Goal Mode Service Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to goal mode service domain categories and the SMF adapter that support this data.

Table A6.11 Tables for the Goal Mode Service Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XTY72GO	RMF GoalModeService	TYPE72GO	Yes	RMF Goal Mode Service Period performance including CPU, memory, I/O, service, response time

The following table shows the domain sub-categories that are relevant to goal mode service domain categories of the SMF adapter.

Table A6.12 Domain Sub-Categories for the Goal Mode Service Domain Category

Domain Category	Domain Sub-Categories
RMF GoalModeService	GoalModeService
	ServicePctResponse
	ServiceExVelocity
	ServiceSysDiscrete

The following table shows the key and ranked metrics that are relevant to goal mode service domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.13 Columns for the Goal Mode Service Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PERFIND	PerformanceIndex	WMean	X	X	Not ranked in ServiceSysDiscrete domain sub-category
		WPStd	X		Used in forecasting **
PCTMISSGOAL	TransMissedGoalPct	WMean	X	X	Not ranked in ServiceSysDiscrete and ServiceExVelocity domain sub-categories
		WPStd	X		Used in forecasting **
PCMETGO	TransMetGoalPct	Min	X		
TSTORE7	TotalStorageUsage	WMean	X	X	
		WPStd	X		Used in forecasting **
PRDMIPS	CpuMipsUsage (Standard Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the PRDMSU column because these columns are ranked in a corresponding manner

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PRDIFAMIPS	IfaCpuMipsUsage (zAAP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the PRDIFAMSU column because these columns are ranked in a corresponding manner
PRDZIPMIPS	ZipCpuMipsUsage (zIIP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the PRDZIPMSU column because these columns are ranked in a corresponding manner
PRDTOTMIPS	AllEngCpuMipsUsage (Software Mips for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the PRDTOTMSU column because these columns are ranked in a corresponding manner
SERVICMSU	TotalServiceMsu	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	
CPUMSU	CpuTcbServiceMsu	Sum	X	X	
IOMSU	IoServiceMsu	Sum	X	X	
SERVICER	TotalServiceRate	WMean	X		
		WPStd	X		Used in forecasting **
RESPAVG	TransAvgResponse	WMean	X	X	

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
		WPStd	X		Used in forecasting **
TRANS	TransCompleted	Sum	X	X	
TRANSRAT	TransCompletedRate	WMean	X		
		WPStd	X		Used in forecasting **
CPUTM	CpuCaptureTime	Sum	X		

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the Workload Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to workload domain categories and the SMF adapter that support this data.

Table A6.14 Tables for the Workload Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XRMFWKL	RMF Workload	RMFWKL RV	Yes	RMF Workload performance using normalized structure including CPU, memory, I/O, service measurements and response time

The following table shows the key and ranked metrics that are relevant to workload domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.15 Columns for the Workload Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
WKLCPPC	CpuBusyPct (Standard Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
WKLIFPC	IfaCpuBusyPct (zAAP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
WKLZPPC	ZipCpuBusyPct (zIIP Engine)	WMean	X	X	

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
		WPStd	X		Used in forecasting **
WKLMEMR	MemoryUsage	Mean	X	X	
		PStd	X		Used in forecasting **
WCLIORT	IoRate	WMean	X	X	
		WPStd	X		Used in forecasting **
WKLEXCP	ExcpRate	WMean	X	X	
		WPStd	X		Used in forecasting **
WKLMIPI	CpuMipsUsage (Standard Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKLMSU column because these columns are ranked in a corresponding manner
WKLIFAMIPS	IfaCpuMipsUsage (zAAP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKLIFAMSU column because these columns are ranked in a corresponding manner
WKLZIPMIPS	ZipCpuMipsUsage (zIIP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKLZIPMSU column because these columns are ranked in a corresponding manner

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
WKLTOTMIPS	AllEngCpuMipsUsage (Software Mips for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKLTOTMSU column because these columns are ranked in a corresponding manner
WKL SMSU	ServiceMsu	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	
WKL SVRT	ServiceRate	WMean	X	X	
		WPStd	X		Used in forecasting **
WKL TRAN	Transactions	Sum	X	X	
WKL RESP	ResponseTime	WMean	X	X	
		WPStd	X		Used in forecasting **

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the Workload Period Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to workload period domain categories and the SMF adapter that support this data.

Table A6.16 Tables for the Workload Period Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XRMFWKP	RMF Workload Period	TYPE72GO RMFWKLRV	Yes	RMF Workload and Workload Period correlated performance measurements created by merging information from MXG data sets TYPE72GO (periods) and RMFWKLRV (workloads)

The following table shows the key and ranked metrics that are relevant to workload period domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.17 Columns for the Workload Period Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PERFIND	PeriodPerformanceIndex	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPCPPC	PeriodCpuBusyPct (Standard Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPIFPC	PeriodIfaCpuBusyPct (zAAP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPZPPC	PeriodZipCpuBusyPct (zIIP Engine)	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPMEMR	PeriodMemoryUsage	Mean	X	X	
		PStd	X		Used in forecasting **
WKPIORT	PeriodIoRate	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPEXCP	PeriodExcpRate	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPMIPS	PeriodCpuMipsUsage (Standard Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKPMSU column because these columns are ranked in a corresponding manner

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
WKPIFAMIPS	PeriodIfaCpuMipsUsage (zAAP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKPIFAMSU column because these columns are ranked in a corresponding manner
WKPZIPMIPS	PeriodZipCpuMipsUsage (zIIP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKPZIPMSU column because these columns are ranked in a corresponding manner
WKPTOTMIPS	PeriodAllEngCpuMipsUsage (Software Mips for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	This ranking can also be used with the WKPTOTMSU column because these columns are ranked in a corresponding manner
WKPSMSU	PeriodServiceMsu	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	
WKPSVRT	PeriodServiceRate	WMean	X	X	
		WPStd	X		Used in forecasting **
WKPTRAN	PeriodTransactions	Sum	X	X	
WKPRESP	PeriodResponseTime	WMean	X	X	

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
		WPStd	X		Used in forecasting **

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the Device Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to device domain categories and the SMF adapter that support this data.

Table A6.18 Tables for the Device Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XTY74	RMF Device	TYPE74	Yes	RMF Device performance

The following table shows the domain sub-categories that are relevant to system device domain category of the SMF adapter.

Table A6.19 Domain Sub-Categories for the Device Domain Category

Domain Category	Domain Sub-Categories
RMF Device	Disk
	Tape

The following table shows the key and ranked metrics that are relevant to device domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.20 Columns for the Device Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
AVENQUE	DeviceQueueAvgLength	WMean	X	X	Not a key metric or ranked metric in Tape domain sub-category
		WPStd	X		Used in forecasting **
PCDVUSE	DeviceBusyPct	WMean	X	X	
		WPStd	X		Used in forecasting **

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PCDVACT	DeviceActivePct	WMean	X		
		WPStd	X		Used in forecasting **
IORATE	DeviceIoRate	WMean	X	X	
		WPStd	X		Used in forecasting **
AVMSRSP	DeviceResponseAvg MSec	WMean	X	X	
		WPStd	X		Used in forecasting **

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the Paging and Swapping Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to paging and swapping domain categories and the SMF adapter that support this data.

Table A6.21 Tables for the Paging and Swapping Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XTY71	RMF PagingAndSwapping	TYPE71	Yes	Older legacy table present in ITRM releases prior to ITRM 3.2. Many metrics are included in the new XRMFSYS super table which contains many system wide metrics.

The following table shows the key and ranked metrics that are relevant to paging and swapping domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.22 Columns for the Paging and Swapping Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PAGRT	PagingRate	WMean	X	X	
		WPStd	X		Used in forecasting **
SWAP	SwapRate	WMean	X	X	
		WPStd	X		Used in forecasting **

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the CPU Processor Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to CPU processor domain categories and the SMF adapter that support this data.

Table A6.23 Tables for the CPU Processor Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XTY70CP	RMF Processor	TYPE70	Yes	RMF CPU processor performance

The following table shows the domain sub-categories that are relevant to CPU processor domain category of the SMF adapter.

Table A6.24 Domain Sub-Categories for the CPU Processor Domain Category

Domain Category	Domain Sub-Categories
RMF Processor	ProcessorCpu
	EngineIfaZip

The following table shows the key and ranked metrics that are relevant to CPU processor domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.25 Columns for the CPU Processor Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
PCCPUBY	CpuBusyPct	WMean		X	Not ranked in EngineIfaZip domain sub-category
PCIZBSY	IfaZipLogicalCpuBusyPct	WMean		X	Not ranked in ProcessorCpu domain sub-category

SMF Adapter Key and Ranked Metrics for the Paging File Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to paging file domain categories and the SMF adapter that support this data.

Table A6.26 Tables for the Paging File Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XTY75	RMF PagingFile	TYPE75	Yes	RMF Paging file performance

The following table shows the key and ranked metrics that are relevant to paging file domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.27 Columns for the Paging File Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
DSBUSY	PagingFileUsagePct	WMean		X	

SMF Adapter Key and Ranked Metrics for the Batch Job, TSO and OMVS Session Performance Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to batch job, TSO, and OMVS session performance domain categories and the SMF adapter that support this data.

Table A6.28 Tables for the Jobs Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XJOBS	SMF Jobs, TSO and OMVS	JOBS	Yes	SMF Batch Job, TSO and OMVS activity and performance

The following table shows the domain sub-categories that are relevant to batch job, TSO, and OMVS session performance domain categories of the SMF adapter.

Table A6.29 Sub-Categories for the Jobs Domain Category

Domain Category	Domain Sub-Categories
SMF Jobs	JobSummary
	JobDetail
	JobAbends
	JobClass
	JobGroup

Domain Category	Domain Sub-Categories
	JobUser
	JobSysplex
	JobMachine
	JobSystemName
	JobServiceClass
	JobTape
SMF TSO	TsoSummary
	TsoDetail
	TsoAbends
	TsoGroup
	TsoUser
	TsoSysplex
	TsoMachine
	TsoSystemName
	TsoServiceClass
SMF OMVS	OmvsSummary
	OmvsDetail
	OmvsAbends
	OmvsGroup
	OmvsUser
	OmvsSysplex
	OmvsMachine
	OmvsSystemName
	OmvsServiceClass

The following table shows the key and ranked metrics that are relevant to jobs domain sub-category and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.30 Columns for the Jobs Domain Sub-Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
JOBCOUNT	JobCount	Sum	X	X	Key and ranked metric only in Job****x domain sub-categories
TSOCOUNT	TsoSessionCount	Sum	X	X	Key and ranked metric only in Tso****x domain sub-categories
OMVSCOUNT	OmvSessionCount	Sum	X	X	Key and ranked metric only in Omvs****x domain sub-categories
JEAPSTM	ElapsedTime	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		
CPUTM	CpuUsageTime	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		
IOMTOTL	IoConnectTime	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		
SYSMIPS	SystemCpuMipsUsage (Standard Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	
IFAMIPS	IfaCpuMipsUsage (zAAP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
ZIPMIPS	ZipCpuMipsUsage (zIIP Engine Software Mips)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X		
TOTMIPS	AllEngCpuMipsUsage (Software Mips for all Engines)	WMean	X		
		WPStd	X		Used in forecasting **
		Sum	X	X	
SERVICMSU	TotalServiceMsu	Sum	X	X	
CPUMSU	CpuTcbServiceMsu	Sum	X	X	
IOMSU	IoServiceMsu	Sum	X	X	
SPOLINE	LinesOutputSpooled	Sum	X	X	Key and ranked metric only in Job****x domain sub-categories
ABENDS	Abends	Sum	X		
TAEDRVS	TapeDrivesAllocated	Sum	X		Key and ranked metric only in Job****x domain sub-categories

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the CICS Transaction Activity Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to CICS transaction activity domain categories and the SMF adapter that support this data.

Table A6.31 Tables for the CICS Transaction Activity Domain Category

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XCICSUM	SMF CICS	ASUMCICX	Yes	SMF CICS activity and performance

The following table shows the domain sub-categories that are relevant to CICS transaction activity domain category of the SMF adapter.

Table A6.32 Domain Sub-Categories for the CICS Transaction Activity Domain Category

Domain Category	Domain Sub-Categories
SMF CICS	CicsEnvironment
	CicsApplid
	CicsApplidTran
	CicsSystem
	CicsTran
	CicsTranApplidUser
	CicsUserTran

The following table shows the key and ranked metrics that are relevant to CICS transaction activity domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.33 Sub-Categories for the CICS Transaction Activity Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
NUMTRAN	TransCount	Sum	X	X	
FCAMCNT	FileAccessCount	Sum	X	X	
TACPUTM	TotalTransCpuTime	Sum	X	X	
IRESPTM	TotalTransResponseTime	Sum	X		
AVGCPU	TransCpuTime	WMean	X		
		WPStd	X		Used in forecasting **
AVGRES	TransResponseTime	WMean	X	X	
		WPStd	X		Used in forecasting **
RSPCNT1	ResponseLessThan1SecondCount	Sum	X		
RESPONSEPCT1	ResponseLessThan1SecondPct	Computed as Percent	X		
RSPCNT2	ResponseLessThan2SecondCount	Sum	X		

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
RESPONSEPCT2	ResponseLessThan2SeondsPct	Computed as Percent	X		
RESPONSECUM PCT2	ResponseLessThan2SeondsCumPct	Computed as Percent	X		
RSPCNT3	ResponseLessThan3SeondCount	Sum	X		
RESPONSEPCT3	ResponseLessThan3SeondsPct	Computed as Percent	X		
RESPONSECUM PCT3	ResponseLessThan3SeondsCumPct	Computed as Percent	X		
RSPCNT4	ResponseLessThan4SeondCount	Sum	X		
RESPONSEPCT4	ResponseLessThan4SeondsPct	Computed as Percent	X		
RESPONSECUM PCT4	ResponseLessThan4SeondsCumPct	Computed as Percent	X		
RSPCNT5	ResponseLessThan5SeondCount	Sum	X		
RESPONSEPCT5	ResponseLessThan5SeondsPct	Computed as Percent	X		
RESPONSECUM PCT5	ResponseLessThan5SeondsCumPct	Computed as Percent	X		
RSPCNT6	ResponseLessThan8SeondCount	Sum	X		
RESPONSEPCT6	ResponseLessThan8SeondsPct	Computed as Percent	X		
RESPONSECUM PCT6	ResponseLessThan8SeondsCumPct	Computed as Percent	X		
RSPCNT7	ResponseLessThan10SecondCount	Sum	X		
RESPONSEPCT7	ResponseLessThan10SecondsPct	Computed as Percent	X		
RESPONSECUM PCT7	ResponseLessThan10SecondsCumPct	Computed as Percent	X		
RSPCNT8	ResponseGreaterThan10SecondCount	Sum	X		

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
RESPONSEPCT7	ResponseGreaterThan 10SecondsPct	Computed as Percent	X		

** To calculate WMean or Mean plus 1 or 2 standard deviations

SMF Adapter Key and Ranked Metrics for the DB2 Thread Accounting Domain Category

The following table shows the ITRM table that contains the key and ranked metrics that are relevant to DB2 thread accounting domain categories and the SMF adapter that support this data.

Table A6.34 *Tables for the DB2 Thread Accounting Domain Category*

ITRM Table	ITRM Table Description	MXG Data Set Name	ASW Available Table	Comments
XDBACCT	SMF DB2	DB2ACCT	Yes	SMF DB2 activity and performance

The following table shows the domain sub-categories that are relevant to DB2 thread accounting activity domain categories of the SMF adapter.

Table A6.35 *Sub-Categories for the DB2 Thread Accounting Domain Category*

Domain Category	Domain Sub-Categories
DB2 Batch	DB2BatchAttach
	DB2BatchJobAttach
	DB2BatchConName
	DB2BatchJobConName
	DB2BatchJob
	DB2BatchLocation
	DB2BatchPlan
	DB2BatchJobPlan
DB2 Cics	DB2CicsConName
	DB2CicsCorrelate
	DB2CicsPlan
	DB2CicsTrans

Domain Category	Domain Sub-Categories
DB2 Ddf	DB2CicsTransPlan
	DB2DdfAttach
	DB2DdfJobAttach
	DB2DdfClientAppl
	DB2DdfNetLocation
	DB2DdfServClass
DB2 Global	DB2Environment
	DB2Subsystem
	DB2Attach
	DB2ConType
	DB2ThreadType
	DB2ConName
	DB2AttachConName
	DB2Correlate
	DB2JobAttach
	DB2Location
	DB2Plan
	DB2Trans

The following table shows the key and ranked metrics that are relevant to DB2 thread accounting domain categories and the SMF adapter that support this data. An X indicates that the key metric or ranked metric listed in the row is supported by the adapter domain category in the corresponding column.

Table A6.36 Columns for the DB2 Thread Accounting Domain Category

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
ACCTCOUNT	AccountingThreadsCount	Sum	X	X	
QWCARNE	IoWaits	Sum	X	X	

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
QWCARNL	LockLatchWaits	Sum	X		
QWACASC	ElapsedTimeInDB2	Sum	X	X	
QWACASCAVG	ElapsedTimeInDb2Avg	Computed as Mean	X		
QWCAJST	CpuTcbTimeInDB2	Sum	X	X	
QWCAJSTAVG	CpuTcbTimeInDb2Avg	Computed as Mean	X		
QWCAWTI	ElapsedIoTimeInDb2	Sum	X	X	
QWCAWTIAVG	ElapsedIoTimeInDb2Avg	Computed as Mean	X		
QWCAWTL	LockLatchTimeInDb2	Sum	X	X	
QWCAWTLAVG	LockLatchTimeInDb2Avg	Computed as Mean	X		
QLCCNVR	ConversationsReceivedFromDdfSite	Sum	X		Key metric only in DB2Ddf****x domain sub-categories
QLCCNVS	ConversationsSentToDdfSite	Sum	X		Key metric only in DB2Ddf****x domain sub-categories
QLCROWR	DataRowReceivedFromDdfSite	Sum	X		Key metric only in DB2Ddf****x domain sub-categories
QLCROWS	DataRowSentToDdfSite	Sum	X		Key metric only in DB2Ddf****x domain sub-categories
QLCCPUR	ElapsedTimeDbAccessAtRemote	Sum	X		Key metric only in DB2Ddf****x domain sub-categories
QLCCPUL	ElapsedTimeLocalWaitingForRemote	Sum	X		Key metric only in DB2Ddf****x domain sub-categories
QLCBYTR	ReceivedFromDdfSiteBytes	Sum	X		Key metric only in DB2Ddf****x domain sub-categories

Column Name	Column Label	Statistic	Key Metric	Ranked Metric	Comments
QLCBYTS	SentToDdfSiteBytes	Sum	X		Key metric only in DB2Ddf****x domain sub-categories

Appendix 7

Data Model

Overview of the SAS IT Resource Management Data Model	459
About the SAS IT Resource Management Data Model	459
Features of the SAS IT Resource Management Data Model	460
Data Model Staging and Aggregation General Concepts	461
Data Model Aggregation Table Groups and Time Periods	462
Data Model Aggregation Key and Ranked Metrics	462
Data Model Aggregation Strategies	463
Data Model Adapter and Domain Category Summary	467
Data Model Adapter and Domain Category Specifications	470
DT Perf Sentry Adapter - DT Perf Sentry Server-Based Domain Categories	470
DT Perf Sentry Adapter - DT Perf Sentry Exchange Domain Category	471
BMC Perf Mgr Adapter - BMC Perf Mgr Domain Category	472
HP Perf Agent and HP Reporter Adapters	474
SAPR3 Adapter - Systems, Machines, Tasks, and Transactions	475
SAPR3 Adapter - SAPBTCH (Batch Job) and SAPWKLD (Workload)	477
SAR Adapter - SAR Domain Category	478
SMF Adapter - RMF Domain Category	479
SMF Adapter - Jobs, TSO, and OMVS Domain Categories	483
SMF Adapter - CICS Domain Category	487
SMF Adapter - DB2 Domain Category	490
SNMP Adapter - Network Activity, Device and Interface Data, and System Data	496
VMware Adapter - Guest and Host Disk, Network, and System Domain Categories	497
MS System Center Operations Manager - MS SCOM Server- Based Domain Categories	498
MS System Center Operations Manager Adapter - MS SCOM Exchange Domain Category	499
Web Log Adapter - All Domain Categories	500

Overview of the SAS IT Resource Management Data Model

About the SAS IT Resource Management Data Model

The data model for SAS IT Resource Management provides a rich set of features. The data model supports a number of data sources from computer system hardware,

operating system software, networks, Web servers, databases, and applications on z/OS, Windows and UNIX platforms. Distinct data sources are characterized as adapters in SAS IT Resource Management terminology. Each distinct adapter can have one or more domain categories that are designed to provide a reporting framework for the adapter data sources. Default domain categories have been created that permit source data to be grouped together for staging and aggregation in order to create report-ready data.

Features of the SAS IT Resource Management Data Model

The following features are available from the SAS IT Resource Management data model for all of the supported adapters:

- Column labeling is standardized to make the data model meaningful and consistent across adapters in SAS IT Resource Management objects, including template tables, staged tables, aggregation tables, and information maps. Labels are coded for easier viewing and to shorten total label length, making them easier to understand. In most cases, columns with similar meaning in different adapters are named with a similar label in order to promote more consistency in classification and metric names across adapters. In most labels, the name incorporates the respective domain category.
- A significant number of computed columns are added to the data model. The function of computed columns is similar to derived and formula variables in the SAS IT Resource Management 2.7 release and in earlier releases. However, one significant difference is that for SAS IT Resource Management, computed columns always take up physical storage. The computations for these computed columns are often simply represented as an RVALUE expression. However, there are some computed columns that include multiple SAS statements in the expression. In addition to these supplied computed columns, users can extend the function of the data model by adding their own computed columns.
- Supplied computed columns include the following areas of functionality:
 - standard date and time derivations that are based on the source datetime found in the raw data.
 - normalized columns that convert lowest common denominator units into industry standard data representations. These normalizations include the following conversions:
 - bytes to kilobytes, megabytes, and gigabytes
 - seconds to milliseconds
 - bits to megabits
 - service units to MSU (millions of service units)
 - percentages between 0 and 1 to percentages between 0 and 100
 - generating a total value from several disparate parts—for example, input + output, read + write, received + sent, and so on.
 - generating a column that contains a value that is the opposite of an already existing column. This conversion highlights the variance in the usage of a resource in contrast to the availability of the same resource.
 - generating event counters to count individual events and enable them to be easily summarized when aggregating the data.
 - generating new classification columns when an individual instance of a performance metric needs to be separated from the global instance. This enables the global and individual instances to be aggregated separately.

- converting raw counts to rate-based values.
- calculation of buckets for counts and percentages that are associated with ranges of response time.
- Formulas, that were present in SAS IT Resource Management 2.7 and earlier, are also available in the new data model, but their usage is different. In SAS IT Resource Management, formulas are reusable shared expressions for use with computed columns. Supplied SAS IT Resource Management formulas are frequently used to set the values for date-based and time-based computed columns that are included in the data model. Customers can also use these same formulas for the computed columns that they create, or they can create new formulas that can be used for their computed columns.
- The most suitable national language support (NLS) format is used for many supplied columns in the data model. Therefore, a large percentage of columns is ready for formatting in multiple languages.

However, some time-based columns that focus on accumulations of time and other columns already using specialized formats have not been converted to NLS because there is currently not an acceptable equivalent NLS format.

Data Model Staging and Aggregation General Concepts

The process of staging data is required in order to prepare raw data and to create data extracts in SAS data set format. Staging is primarily an intermediate step that is performed before the creation of aggregations in the SAS IT Resource Management data model. (Existing staged data is overwritten for each new iteration of the staging process.) The SAS IT Resource Management data model is based on aggregations. Aggregations are created to provide report-ready data that can be easily input into an information map. An information map is a map that enables the aggregated data to be available for several SAS reporting clients, including SAS Enterprise Guide, SAS Web Report Studio, and SAS OLAP Cube Studio.

Note: The following information describes the aggregations that are generated by means of the Adapter Setup wizard. The information does not necessarily apply to all aggregations in general.

Aggregated data tables in the SAS IT Resource Management data model have the following structures and properties associated with them:

- The name of the aggregation table consists of a concatenation of a time period and a descriptive string. The descriptive string represents a domain category in a mixed-case format that uses uppercase letters to signify the beginning of a new word. For example, the name of an aggregation table might be one of the following: DayHourCpu or MonthJobSummary.
- The aggregation table has an aging limit specified in days.
- The time period for aggregated data is represented by one or more columns in the classification list.
- The aggregated data contains one or more classification columns. Some of these columns are time-based and other columns are not based on time.
- The aggregated data might contain one or more ID columns.
- The aggregated data contains one or more statistics. Statistics can be weighted or unweighted, depending on the nature of the performance metric. Weighting is

typically by duration of time or by a counter that indicates a number of events. Weighting columns are explicitly specified in SAS IT Resource Management.

- The aggregated data can contain one or more ranked metrics. Each ranked metric can be based on a class or ID variable or on a statistic.
- The aggregated data can contain one or more computed columns. These columns can use any aggregated data as sources for their calculations.
- Filters can be used both to keep only appropriate data, and to reduce the volume of the aggregation table's output data.

Data Model Aggregation Table Groups and Time Periods

Aggregation tables in the data model are grouped into aggregation table groups. Multiple aggregation table groups can be defined within an adapter's domain category. An aggregation table group is a set of aggregation tables that contain a set of classification columns that are identical, except for the time period classification columns. The time periods vary among the aggregation tables in the aggregation table group.

For example, in an aggregation table group that is based on Memory, the aggregation tables are named DayMemory, DayHourMemory, DayShiftMemory, MonthMemory, MonthHourMemory, MonthShiftMemory, and XMinMemory. (XMinMemory signifies aggregation tables that are based on sub-hourly memory activity). Typical time periods that can be represented in an aggregation table are datetime, hour, shift, day, week, and month. An aggregation table can focus on a single time period, such as datetime, day, week, or month. Alternatively, it can include multiple time periods such as day and hour or day and shift. An aggregation table that uses the datetime period typically focuses on time intervals that are less than a full hour. Aggregation tables that focus on day, week, month, hour, and shift time periods work with time periods of an hour or more.

Data Model Aggregation Key and Ranked Metrics

The SAS IT Resource Management data model provides metrics that are identified as key or ranked metrics. Key and ranked metrics are primarily the same metrics within the adapters, but they are handled differently depending on the type of aggregation that they are used in.

- Key metrics are metrics that are output to a Key Metrics aggregation table. A Key Metrics aggregation table usually includes the term “KeyMetrics” in its name (for example, KeyMetricsMemory). The intent of the Key Metrics aggregation table is to keep a limited set of important performance metrics for a long period of time. This tactic facilitates better capacity planning and forecast reporting. Key metrics are not ranked. Only one statistic is specified for a key metric in a Key Metrics aggregation table.
- Ranked metrics are metrics within an aggregation table for which ranking is done. Metrics can be ranked in any aggregation tables, but they are typically not ranked in Key Metrics or XMin (sub-hourly) aggregation tables. Ranked metrics are based on statistics. Typically only the most meaningful statistic is ranked.
- Both metrics and date classifications can be ranked using either ascending or descending criteria that is based on a list of classifications that are specified for the ranking of the available data.

Data Model Aggregation Strategies

There are several distinct types of aggregation processing strategies that are available for the supported adapters and domain categories in the data model. These strategies are listed here with characteristics that are representative of them.

- Server activity strategies are used for the following adapters and domain categories:
 - RMF domain category of the SMF adapter
 - DT Perf Sentry Exchange InfoStore domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Exchange InfoStoreMailbox domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Exchange InfoStorePrivate domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Exchange InfoStorePublic domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Exchange MailServer domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Server Disk domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Server Network domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Server PagingFile domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Server PhysicalDisk domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Server ProcessorCpu domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - DT Perf Sentry Server System domain category of the DT Perf Sentry adapter (formerly the NTSMF adapter)
 - HP Perf Agent Disk UNIX domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
 - HP Perf Agent Disk Windows domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
 - HP Perf Agent LogicalSystem UNIX domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
 - HP Perf Agent Network UNIX domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
 - HP Perf Agent Network Windows domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
 - HP Perf Agent ProcessorCpu UNIX domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
 - HP Perf Agent ProcessorCpu Windows domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)

- HP Perf Agent System UNIX domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
- HP Perf Agent System Windows domain category of the HP Perf Agent adapter (formerly the HP OpenView Performance Agent, or HP-OVPA, adapter)
- HP Reporter Disk UNIX domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter Disk Windows domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter LogicalSystem UNIX domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter Network UNIX domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter Network Windows domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter ProcessorCpu UNIX domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter ProcessorCpu Windows domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter System UNIX domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- HP Reporter System Windows domain category of the HP Reporter adapter (formerly the HP OpenView Reporter, or HP-OVREP, adapter)
- BMC Perf Mgr CacheWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr CpuUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr CpuWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr DiskUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr FilesystemUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr KernelUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr LogicalDiskWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr MemoryUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr MemoryWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr NetworkUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr NetworkWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)

- BMC Perf Mgr PagingFileWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr PhysicalDiskWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr ProcessUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr SwapUnix domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- BMC Perf Mgr SystemWindows domain categories of the BMC Perf Mgr adapter (formerly the BMC Patrol adapter)
- SAR domain category of the SAR adapter
 - Processor CPU
 - Network Activity
 - System, Global CPU, Queuing, Memory, Paging, Swap, Cache, Kernel
- SCOM Server Disk domain category of the MS SCOM adapter
- SCOM Server Network domain category of the SMS SCOM adapter
- SCOM Server Paging File domain category of the MS SCOM adapter
- SCOM Server PhysicalDisk domain category of the MS SCOM adapter
- SCOM Server ProcessorCpu domain category of the MS SCOM adapter
- SCOM Server System domain category of the MS SCOM adapter
- SCOM Exchange Disk domain category of the MS SCOM adapter
- SCOM Exchange InfoStore domain category of the MS SCOM adapter
- SCOM Exchange InfoStoreMailbox domain category of the MS SCOM adapter
- SCOM Exchange InfoStorePublic domain category of the MS SCOM adapter
- SCOM Exchange MailServer domain category of the MS SCOM adapter
- SCOM Exchange Network domain category of the MS SCOM adapter
- SCOM Exchange Paging File domain category of the MS SCOM adapter
- SCOM Exchange PhysicalDisk domain category of the MS SCOM adapter
- VMware Guest Disk domain category of the VMware adapter
- VMware Guest Network domain category of the VMware adapter
- VMware Guest System domain category of the VMware adapter
- VMware Host Disk domain category of the VMware adapter
- VMware Host Network domain category of the VMware adapter
- VMware Host System domain category of the VMware adapter

The server activity aggregation strategies are identified by the following characteristics:

- Data collection occurs at designated intervals.
- Server activity in various domain categories (CPU, Memory, Paging, Disk, Network, and so on) is the primary focus.
- There is one aggregation table group per Aggregation transformation.

- In general, there are only a few classification columns in the aggregation tables.
- There are usually diverse metrics in the various aggregation table groups that are created.
- Statistics are mostly explicitly weighted by the DURATION column. Common statistics are mean, weighted mean, maximum, minimum, weighted population standard deviation, and population standard deviation.
- Aggregation tables are typically created for eleven time periods. Because data volume is typically low to moderate, it is useful to see server activity across many time periods.
- Supported domain categories focus on identifying TOPN and BOTTOMN instances for the ranked metrics in these aggregation tables and reporting on these instances.
- Event activity strategies are used for the following adapters and domain categories:
 - Jobs (batch jobs) domain category of the SMF adapter
 - TSO (sessions) domain category of the SMF adapter
 - DB2 (accounting threads) domain category of the SMF adapter
 - Batch Job and Workload domain categories of the SAPR3 adapter

The event activity aggregation strategies are identified by the following characteristics:

- Data collection occurs at termination of an event.
- Event activity (jobs, tasks, transactions, threads, sessions, other event types) are the focus.
- Commonly the same list of metrics is found among many aggregation table groups. The list of classifications is what typically varies here.
- With more classifications, many aggregation table groups are typically created from a single staged table. The goal is to create smaller and more compact aggregation tables that are based on subsets of classifications available in the full classification list. Working with these smaller aggregations should significantly speed up reporting.
- Larger lists of classification columns are typically found in this strategy.
- Statistics, when weighted, use mostly a count-based column (for example, a count of jobs, transactions, or samples). The most common statistics specified are sum, weighted mean, and maximum.
- Because many events are typically collected, data volume is typically high to extremely high.
- Because of this high volume, aggregation tables are typically created for fewer time periods. There are typically no aggregation tables created for the XMin (sub-hourly) time period.
- This strategy uses the following types of metrics for reporting:
 - counts of events
 - response time
 - levels of service
 - totals of performance resources or service units used
 - treatment of the favored events (for example, jobs or transactions)

- identification of events that use excessive resources
- Hourly summarization of event activity strategies is used for the following adapters and domain categories:
 - CICS (transactions) domain category of the SMF adapter
 - SAPSYS (systems), SAPHST (application servers), SAPSMT (task goals), SAPTSK (tasks), and SAPTRN (transactions) domain categories of the SAPR3 adapter

This aggregation strategy is identified by the following characteristics:

- It is not a distinct strategy, but a variation of the event activity strategy.
- The staged table that is used as input to building aggregation tables has already been summarized to the hour. This characteristic distinguishes it from the event activity strategy. In addition, the volume of staged data to be read as input is minimized. It might (or might not) significantly affect the volume in the aggregation output.

Data Model Adapter and Domain Category Summary

Each SAS IT Resource Management supplied adapter can have one or more domain categories that are associated with it. Domain categories have been created that enable users to group source data that should have data staged and aggregated together for reporting purposes. The current mapping of the SAS IT Resource Management data model supplied adapters and domain categories is as follows:

- DT Perf Sentry (eleven domain categories)
 - Server Disk
 - Server Network
 - Server PagingFile
 - Server PhysicalDisk
 - Server ProcessorCpu
 - Server System
 - Exchange InfoStore
 - Exchange InfoStoreMailbox
 - Exchange InfoStorePrivate
 - Exchange InfoStorePublic
 - Exchange MailServer
- BMC Perf Mgr (sixteen domain categories)
 - CacheWindows
 - CpuUnix
 - CpuWindows
 - DiskUnix
 - FilesystemUnix
 - KernelUnix
 - LogicalDiskWindows

- MemoryUnix
- MemoryWindows
- NetworkUnix
- NetworkWindows
- PagingFileWindows
- PhysicalDiskWindows
- ProcessUnix
- SwapUnix
- SystemWindows
- HP Perf Agent (nine domain categories that contain both UNIX and Windows platforms)
 - Disk UNIX
 - Disk Windows
 - LogicalSystem UNIX
 - Network UNIX
 - Network Windows
 - ProcessorCpu UNIX
 - ProcessorCpu Windows
 - System UNIX
 - System Windows
- HP Reporter (nine domain categories that contain both UNIX and Windows platforms)
 - Disk UNIX
 - Disk Windows
 - LogicalSystem UNIX
 - Network UNIX
 - Network Windows
 - ProcessorCpu UNIX
 - ProcessorCpu Windows
 - System UNIX
 - System Windows
- SAR (four domain categories)
 - Processor CPU
 - Network Activity
 - System, Global CPU, Queuing, Memory, Paging, Swap, Cache, Kernel
 - Disk Activity
- SAPR3 (seven domain categories)
 - SAP R/3 System (system measurement)

- SAP R/3 Host (application server measurement)
- SAP R/3 System Host TaskType (task activity by application server)
- SAP R/3 TaskType (task activity by system)
- SAP R/3 Tcode (transaction measurement)
- SAP R/3 Batch (batch job measurement)
- SAP R/3 Workload (workload measurement)
- SMF (four domain categories)
 - RMF
 - Jobs and TSO
 - CICS
 - DB2
- SNMP (six domain categories)
 - Device ad Interface data
 - System data
 - IP protocol data
 - TCP protocol data
 - ICMP protocol data
 - UDP protocol data
- Microsoft System Center Operations Manger (fourteen domain categories)
 - Server Disk domain category of the MS SCOM adapter
 - Server Network domain category of the SMS SCOM adapter
 - Server Paging File domain category of the MS SCOM adapter
 - Server PhysicalDisk domain category of the MS SCOM adapter
 - Server ProcessorCpu domain category of the MS SCOM adapter
 - Server System domain category of the MS SCOM adapter
 - Exchange Disk domain category of the MS SCOM adapter
 - Exchange InfoStore domain category of the MS SCOM adapter
 - Exchange InfoStoreMailbox domain category of the MS SCOM adapter
 - Exchange InfoStorePublic domain category of the MS SCOM adapter
 - Exchange MailServer domain category of the MS SCOM adapter
 - Exchange Network domain category of the MS SCOM adapter
 - Exchange Paging File domain category of the MS SCOM adapter
 - Exchange PhysicalDisk domain category of the MS SCOM adapter
- VMware vCenter (eight domain categories)
 - Guest Disk
 - Guest Network
 - Guest Storage

- Guest System
- Host Disk
- Host Network
- Host Storage
- Host System
- Web Log (four domain categories)
 - Browser Activity
 - Server Activity
 - Visitor Organization Activity
 - Web Page Activity

Note: Other adapters are supported in the Adapter Setup wizard but have no supplied aggregations. These adapters include DCOLLECT, EREP, IMF, NTSMF-MXG, TMON2CIC, TMONDB2, TMS, TPF, and VMMON.

Data Model Adapter and Domain Category Specifications

DT Perf Sentry Adapter - DT Perf Sentry Server-Based Domain Categories

These domain categories represent a subset of data collected by Performance Sentry software that is available from Demand Technology Software. They focus on server-level activity and performance on the Windows platform. These adapter and domain category combinations use the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.1 Domain Categories for Staged Tables (DT Perf Sentry Server Domain Category)

Domain Category	Staged Table
Server ProcessorCpu	NTPRCSR
Server Network	NTINTRF
Server System	NTSYSTEM_ITRM
Server Disk	NTLGDSK
Server PhysicalDisk	NTPHDSK
Server PagingFile	NTPGNFL

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables.

The following key classifications are found in this domain category:

- Domain
- System (Machine)
- CPU ID
- Disk ID
- Network Interface

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- utilization of CPU, memory, queuing, cache, and paging on the global server level
- utilization of CPU for individual CPU processors
- utilization, response times, transfer rates, and wait and queuing for logical and physical disks
- packet and data volume transfer activity, utilization, and queuing for network interfaces

DT Perf Sentry Adapter - DT Perf Sentry Exchange Domain Category

This domain category represents a subset of data that is collected by Performance Sentry software, which focuses on Microsoft Exchange server activity and performance on the Windows platform. This adapter and domain category combination uses the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.2 *Domain Categories for Staged Tables (DT Perf Sentry Exchange Domain Category)*

Domain Category	Staged Table
Exchange Mailbox Information Store	NTSMLBX
Exchange Information Store	NTMSEIS
Exchange Public Information Store	NTSPBLC
Exchange Private Information Store	NTSPRVT

Domain Category	Staged Table
Mail Server Activity	NTMSRVR

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables.

The following key classifications are found in this domain category:

- Domain
- System (Machine)
- Information Store ID
- Mail Server ID

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- Remote Procedure Call (RPC) response times, transfer rates, and request and operation rates at the global Exchange server level
- total and active user counts at the global Exchange server level
- virtual memory utilization at the global Exchange server level
- delivered, sent, and submitted messages and receive and send queue lengths for mailboxes and public and private information stores
- message rates, data transfer rates, and queue lengths for Exchange mail servers

BMC Perf Mgr Adapter - BMC Perf Mgr Domain Category

This adapter and domain category combination uses the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.3 Domain Categories for Staged Tables (BMC Perf Mgr Domain Category)

Domain Category	Staged Table
CacheWindows	PTCACHE
CpuUnix	PTUCPU
CpuWindows	PTNTCPU

Domain Category	Staged Table
DiskUnix	PTUDISK
FilesystemUnix	PTUFSTM
KernelUnix	PTUKRNL
LogicalDiskWindows	PTLGDSK
MemoryUnix	PTUMMRY
MemoryWindows	PTMEMRY
NetworkUnix	PTUNWRK
NetworkWindows	PTNTWRK
Paging FileWindows	PTPGFLE
PhysicalDiskWindows	PTPHDSK
ProcessUnix	PTUPRCS
SwapUnix	PTUSWAP
SystemWindows	PTNTSYS

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables. Filters that query the CPU instance are used to divide aggregation tables that contain global CPU measurements from those that contain individual CPU processor measurements.

The following key classifications are found in this domain category:

- System (Machine)
- CPU ID
- Disk ID
- Filesystem ID
- Network Interface

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- utilization of CPU, memory, queuing, cache, paging, and swapping on the global server level

- utilization of CPU for individual CPU processors
- utilization, response times, transfer rates, and wait and queuing for physical disks
- utilization of space for file systems
- packet and data volume transfer activity, utilization, and queuing for network interfaces

HP Perf Agent and HP Reporter Adapters

This group of adapter and domain categories consists of the following entities that apply to both the HP Perf Agent and HP Reporter adapters:

- Disk UNIX
- Disk Windows
- LogicalSystem UNIX
- Network UNIX
- Network Windows
- ProcessorCpu UNIX
- ProcessorCpu Windows
- System UNIX
- System Windows

These domain categories focus on server level activity and performance on the UNIX and Windows platforms. With recent changes made to data collection by HP Perf Agent and HP Reporter, it is possible to monitor the performance of logical systems that are configured as virtual guests of host servers. These adapters provide some support of the new technology of virtualization. Operating environments that are currently supported for virtualization are HP Virtual machines, AIX LPARs, and VMWARE ESX guests on Linux. These adapters use the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.4 Domain Categories for Staged Tables (HP Perf Agent and HP Reporter Domain Categories)

Domain Category	Staged Table
Processor CPU	PCSCPU
Network Activity	PCSNET
System, Global CPU, Queuing, Memory, Paging, Swap, Cache, Kernel, Global Disk, Global Network	PCSGLB
Disk Activity	PCSDSK
Logical and Virtual System CPU and Memory	PCSLS

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables. Filters that query the operating system are used to divide aggregation tables that contain UNIX platform performance measurements from those that contain Windows platform performance measurements.

The following key classifications are found in this domain category:

- System (Machine)
- CPU ID
- Disk ID
- Network Interface
- Logical System ID
- Logical System name

In all the domain categories listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- utilization of CPU, memory, queuing, paging, swap, disk, and network on the global server level
- utilization of CPU for individual CPU processors
- utilization, response times, transfer rates, and wait and queuing for physical disks
- packet and data volume transfer activity, utilization, queuing, and error rates for network interfaces
- utilization of CPU and memory for logical systems to determine the share of resource utilization for the virtual guests associated with host servers

SAPR3 Adapter - Systems, Machines, Tasks, and Transactions

This group of adapter and domain categories consists of the following entities:

- SAPR3 Adapter - SAP R/3 System Domain Category
- SAPR3 Adapter - SAP R/3 Host Domain Category
- SAPR3 Adapter - SAP R/3 System Host TaskType Domain Category
- SAPR3 Adapter - SAP R/3 System TaskType Domain Category
- SAPR3 Adapter - SAP R/3 TransactionCode Domain Category

SAP R/3 performance data is collected by using the SAS IT Management Adapter for SAP in conjunction with SAS/ACCESS to SAP R/3 software. This adapter and domain category combination uses the event hourly summarization strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged

tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.5 Domain Categories for Staged Tables (SAPR3 Adapter Domain Categories for Systems, Machines, Tasks, and Transactions)

Domain Category	Staged Table
System Summary	SAPSYS
System-Host Summary	SAPHST
System-Host-TaskType Summary	SAPSMT
System-TaskType Summary	SAPTSK
System-Transaction	SAPTRN

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, and monthly time periods. Hourly activity is included in a number of these aggregation tables. (There are no sub-hourly aggregation tables for this domain category due to the staged data being pre-summarized to the hourly level.) There are groups of aggregation tables that are created for each of the staged tables in the previous table. Filters that query the transaction code are used to exclude blank transaction codes from the system-transaction aggregation tables.

The following key classifications are found in this domain category:

- System
- Machine (Host or Application Server)
- Account ID
- Task Type
- Transaction Code
- Elapsed Time Goal
- Program Name

There are metrics that are designated as key or ranked metrics. These metrics are included in a special Key Metrics aggregation table that is intended for longer aging periods and used for capacity planning and forecasting reporting. These metrics are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among various classification values within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas for various groups of classifications:

- event, task type, and transaction counts
- total and average values for CPU, database, elapsed, response, processing, overhead, wait, and RFC times
- total and average values for database requests and physical database accesses and changes
- total and average values for kilobytes of data transferred

- request counts and time expended for sequential reads, direct reads, update, insert, and delete operations
- counts of physical records read, updated, inserted, and deleted
- lock counts

SAPR3 Adapter - SAPBTCH (Batch Job) and SAPWKLD (Workload)

SAP R/3 performance data is collected by using the SAS IT Management Adapter for SAP in conjunction with SAS/ACCESS to SAP R/3 software. This adapter and domain categories combination uses the event activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.6 Domain Categories for Staged Tables (SAPBTCH and SAPWKLD Domain Categories)

Domain Category	Staged Table
Batch Jobs	SAPBTCH
Workloads	SAPWKLD

These staged tables are input to a series of aggregation jobs that create aggregation tables for daily, weekly, and monthly time periods. (At this time, there is no hourly, sub-hourly, or shift-based activity in these aggregations.) There are several groups of aggregation tables that are created for each of the staged tables that are listed in the previous table. Filters that query transaction code are used to exclude blank transaction codes from the workload-related aggregation tables that focus on transaction activity.

The following key classifications are found in these domain categories:

- System
- Machine (Host or Application Server)
- Account ID
- Job Name
- Transaction Code
- Module
- Sub-Application
- Development Class
- Task Type
- Program Name

There are metrics that are designated as key or ranked metrics. These metrics are included in a special Key Metrics aggregation table that is intended for longer aging periods and used for capacity planning and forecasting reporting. These metrics are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among various classification values within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?”](#) on page 417.

Report-ready data for this adapter focuses primarily on the following areas for various groups of classifications:

- event, job, and transaction counts
- total and average values for CPU, database, elapsed, response, processing, overhead, wait, and RFC times
- total and average values for database requests and physical database accesses and changes
- total and average values for kilobytes of data transferred
- request counts and time expended for sequential reads, direct reads, update, insert, and delete operations
- counts of physical records read, updated, inserted, and deleted
- lock counts

SAR Adapter - SAR Domain Category

This domain category represents a subset of data collected by the System Activity Reporter (SAR) utility on various UNIX environments and focuses on server level activity and performance on the UNIX platform. This adapter and domain category combination uses the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.7 Domain Categories for Staged Tables (SAR Domain Category)

Domain Category	Staged Table
Processor CPU	SARCPUB
Network Activity	SARIFAC
System, Global CPU, Queuing, Memory, Paging, Swap, Cache, Kernel	SAR
Disk Activity	SARDEV

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables.

The following key classifications are found in this domain category:

- System (Machine)
- CPU ID
- Disk ID
- Network Interface

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity

planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- utilization of CPU, memory, queuing, cache, paging, and swapping on the global server level
- utilization of CPU for individual CPU processors
- utilization, response times, transfer rates, and wait and queuing for physical disks
- packet and data volume transfer activity for network interfaces

SMF Adapter - RMF Domain Category

Resource Management Facility (RMF) data collection represents a subset of the following record types: 70, 71, 72, 74, 75, 77, and 78. These record types are collected by SMF on the z/OS platform. This adapter and domain category combination uses the server activity strategy for staging and aggregating data. MXG software is the primary tool that is used in staging the data. A SAS IT Resource Management staging transformation uses the MXG staged output to complete the staging of the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

The Adapter Setup Wizard Default Availability legend is applicable to the following RMF based tables:

- 1 — Available when the Minimal configuration level is selected
- 2 — Available when the Typical configuration level is selected
- 3 — Available when the Full configuration level is selected

Table A7.8 Domain Categories and Staged Tables Output from the Adapter Staging Job (RMF Report Group)

Domain Category	Staged Table	MXG Feeder Data Sets	Adapter Setup Wizard Default Availability	Comments and Significant Changes for SAS IT Resource Management 3.2
System-wide CPU, Memory, I/O, Queueing, Service Usage, and Paging and Swapping	XRMFSYS	RMFINTRV TYPE71	Yes - 1, 2, 3	New table. New metrics for hardware-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included. Paging and swapping metrics are included.
System-wide CPU, Memory, I/O, Queueing, Workload measurements	XRMFINT	RMFINTRV	No	Legacy table with imbedded Workload measurements that are not normalized. New metrics for hardware-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.

Domain Category	Staged Table	MXG Feeder Data Sets	Adapter Setup Wizard Default Availability	Comments and Significant Changes for SAS IT Resource Management 3.2
Memory, Paging, Swapping	XTY71	TYPE71	Yes - 3	Legacy table. Large number of metrics are included in new XRMFSYS table.
LPAR CPU and Service usage activity by Sysplex per System	XASM70L	ASUM70LP	Yes - 1, 2, 3	New metrics for hardware-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
LPAR CPU and Service usage activity by CEC	XASMCEL	ASUMCELP	Yes - 1, 2, 3	New metrics for hardware-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
Workload Resource utilization	XRMFWKL	RMFWKLRV	Yes - 1, 2, 3	New table. Metrics are normalized by Workload ID and Description to enable easier reporting by workload. New metrics for software-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
Goal Mode Period Resource utilization	XTY72GO	TYPE72GO	Yes - 1, 2, 3	New metrics for software-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
Workload And Workload Period Cross Reference Resource utilization	XRMFWKP	TYPE72GO RMFWKLRV	Yes - 3	New table. Provides side-by-side view of Workloads and the Workload periods that comprise the individual workloads. New metrics for software-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
Processor CPU utilization	XTY70CP	TYPE70	Yes - 3	zAAP and zIIP specialty engine metrics are now included.
Device Activity (Disk and Tape)	XTY74	TYPE74	Yes - 2	No changes.
Paging Data Set Activity	XTY75	TYPE75	Yes - 3	No changes.

These staged tables are input to a number of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in many of these aggregations. These aggregations are structured by domain categories. In some cases, domain categories are further broken down into subsets of aggregations known as sub-categories. The following table summarizes these domain categories and sub-categories.

Table A7.9 Domain Categories for Staged Tables (RMF Report Group)

Domain Category Area	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
System-wide CPU, Memory, I/O, Service usage, Paging, and Swapping	XRMFSYS	RMF System		Yes - 1, 2, 3	
System-wide CPU, Memory, I/O, Queuing, Workload measurements	XRMFINT	RMF Interval		No	
Memory, Paging, Swapping	XTY71	RMF Paging and Swapping		Yes - 3	
LPAR CPU and Service usage activity by Sysplex per System	XASM70L	RMF Lpar		Yes - 1, 2, 3	
LPAR CPU and Service usage activity	XASMCEL	RMF CecLpar		Yes - 1, 2, 3	
Workload Resource utilization	XRMFWKL	RMF Workload		Yes - 1, 2, 3	
Goal Mode Period Resource utilization	XTY72GO	RMF GoalModeService	GoalModeService	Yes - 1, 2, 3	No filters used, so all Response time types are included
Goal Mode Period Resource utilization	XTY72GO	RMF GoalModeService	ServicePctResponse	No	Filter used to select only Percentile Response response type activity
Goal Mode Period utilization	XTY72GO	RMF GoalModeService	ServiceExVelocity	No	Filter used to select only Execution Velocity response type activity
Goal Mode Period utilization	XTY72GO	RMF GoalModeService	ServiceSysDiscrete	No	Filter used to select only System Discretionary response type activity

Domain Category Area	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
Workload and Workload Period Cross Reference Resource utilization	XRMFWKP	RMF WorkloadPeriod		Yes - 3	These aggregations contain a composite set of metrics that cross-reference workloads and workload periods
Processor CPU utilization	XTY70CP	RMF Processor	ProcessorCPU	Yes - 3	Filter used to select only Standard engine processor activity
Processor CPU utilization	XTY70CP	RMF Processor	EnginIfaZip	Yes - 3	Filter used to select only specialty (zAAP and zIIP) engine processor activity
Device Activity (Disk)	XTY74	RMF Device	Disk	Yes - 2, 3	Filter used to select only Disk-related activity
Device Activity (Tape)	XTY74	RMF Device	Tape	Yes - 3	Filter used to select only Tape related activity
Paging Data Set Activity	XTY75	RMF PagingFile		Yes - 3	

The following key classifications are found in this domain category:

- Sysplex
- CEC
- System (Machine)
- CPU ID
- LPAR name
- LPAR number-name concatenation
- Service Class
- Service Class Period
- Workload ID
- Workload ID-Description concatenation
- Workload Period
- Device Number or Volume
- Paging data set name

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity

planning and forecasting reporting. They are also included in other aggregation tables to enable the comparison of resource usage among various classification values within the aggregated data. They also provide the ability to filter on the top or bottom rated resource values within a set of classification values. The latter ranking filters are very useful to limit the number of reports that are created to those classifications with the greatest or the least resource utilization. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following domain categories:

- utilization of CPU, memory, queuing, paging, and swap activity on the global server level
- CPU-based hardware MSU and MIPS usage on the global server level and by Processor engine type
- utilization of CPU and level of service measurements by logical partition (LPAR)
- CPU-based hardware MSU and MIPS usage by LPAR and by Processor engine type
- utilization of CPU for individual CPU processors
- utilization of CPU, memory, I/O, and level of service measurements by defined workload
- CPU-based software MSU and MIPS usage by defined workload and by Processor engine type
- utilization of CPU, memory, I/O, and level of service measurements by defined Service class and period
- CPU-based software MSU and MIPS usage by defined Service class and period and by Processor engine type
- utilization, response times, and wait and queuing for devices (disk and tape)

SMF Adapter - Jobs, TSO, and OMVS Domain Categories

Batch Job, TSO session, and OMVS session data collection represents a subset of record types (6, 26, and 30) collected by SMF on the z/OS platform. The full complement of data that is collected by SMF for these record types contains performance measurements for batch jobs, TSO sessions, started tasks, Open Edition MVS sessions, and APPC tasks. For SAS IT Resource Management 3.2, this domain category focuses solely on Batch job, TSO sessions, and OMVS session activity. This adapter and domain category combination uses the event activity strategy for the staging and aggregating of data. MXG software is the primary tool that is used in staging the data. A SAS IT Resource Management staging transformation uses the MXG staged output as input and then completes the staging of the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.10 Domain Categories and Staged Tables Output from the Adapter Staging Job (Jobs, TSO, and OMVS Domain Categories)

Domain Category	Staged Table	MXG Feeder Data Sets	Adapter Setup Wizard Default Availability	Comments and Significant Changes for SAS IT Resource Management 3.2
Batch Job activity and performance	XJOBS	JOBS	Yes - 1, 2, 3	New metrics for software-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
TSO Session activity and performance	XJOBS	JOBS	Yes - 1, 2, 3	New metrics for software-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included.
OMVS Session activity and performance	XJOBS	JOBS	Yes - 1, 2, 3	New metrics for software-based MSU and MIPS are included. zAAP and zIIP specialty engine metrics are included. OMVS activity performance is new.

This staged table is input to a number of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, and monthly time periods. Shift-based activity is included in a number of these aggregation tables. These aggregations are structured by domain categories. All domain categories in this group are further broken down into subsets of aggregations known as sub-categories. The following table summarizes these domain categories and sub-categories.

Table A7.11 Domain Categories for Staged Tables (Jobs, TSO, and OMVS Domain Categories)

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
Batch Job activity and performance	XJOBS	JOBS	JobAbends	Yes - 1, 2, 3	Filter used to select only Jobs that have abended
			JobClass	Yes - 1, 2, 3	Job activity is summarized by Job class
			JobDetail	Yes - 3	Filter used to select activity for individual jobs that use some resources. These aggregations can be quite high in volume.

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
			JobGroup	Yes - 2, 3	Job activity is summarized by Groups of users
			JobMachine	Yes - 2, 3	Job activity is summarized by Machine
			JobService Class	Yes - 3	Job is summarized by Service class
			JobSummary	Yes - 1, 2, 3	Job activity is summarized by Job name
			JobUser	Yes - 2,3	Job activity is summarized by User
			JobSysplex	No	Job activity is summarized by Sysplex
			JobSystem Name	No	Job activity is summarized by SystemName
			JobTape	No	Filter used to select only Jobs that have used Tape resource
TSO Sessions activity and performance	XJOBS	TSO	TsoAbends	No	Filter used to select only TSO sessions that have abended
			TsoDetail	Yes - 3	Filter used to select activity for individual TSO sessions that use some resources
			TsoGroup	Yes - 2, 3	TSO session is summarized by Groups of users
			TsoMachine	Yes - 2, 3	TSO session activity is summarized by Machine
			TsoService Class	No	TSO session activity is summarized by Service class

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
			TsoSummary	Yes - 2, 3	TSO session activity is summarized by Session name
			TsoSysplex	No	TSO session is summarized by Sysplex
			TsoSystemName	No	TSO session is summarized by System name
			TsoUser	Yes - 2, 3	TSO session activity is summarized by User
OMVS Session activity and performance	XJOBS	OMVS	OmvsAbends	No	Filter used to select only OMVS sessions that have abended
			OmvsDetail	Yes - 3	Filter used to select activity for individual OMVS sessions that use some resources
			OmvsGroup	Yes - 2, 3	OMVS session activity is summarized by Groups of users
			OmvsMachine	Yes - 2, 3	OMVS session activity is summarized
			OmvsServiceClass	No	OMVS session activity is summarized by Service class
			OmvsSummary	Yes - 2, 3	OMVS session activity is summarized by Session name
			OmvsSysplex	No	OMVS session activity is summarized by Sysplex
			OmvsSystemName	No	OMVS session activity is summarized by System name
			OmvsUser	Yes - 2, 3	OMVS session activity is summarized by User

The following key classifications are found in this domain category:

- Sysplex
- System (Machine)
- Job class
- Job name
- Job number, TSO session number, or OMVS session number
- Service class
- Group ID
- User ID

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. Key metrics and specified statistics are included in special Key Metrics aggregations that are intended for longer aging periods and used for capacity planning and forecast reporting. Ranked metrics and specified statistics are included in other aggregation tables to enable the comparison of resource usage among various classification values within the aggregated data. They are also used to filter on the top or bottom rated resource values within a set of classification values. The latter ranking filters are very useful to limit the number of reports that are created to those classifications with the greatest or the least resource utilization. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following domain categories:

- job counts, TSO session counts, and OMVS session counts for various groups of classifications
- accumulated service and times for CPU, memory, and I/O for individual batch jobs, TSO sessions, OMVS sessions across various groups of classifications
- accumulated CPU-based software MSU and MIPS by Engine type for individual batch jobs, TSO sessions, and OMVS sessions across various groups of classifications
- printing output associated with batch jobs and TSO sessions
- utilization of tape resources for batch jobs
- batch job, TSO session, and OMVS abends and associated return codes

SMF Adapter - CICS Domain Category

CICS data represents record type 110 collected by SMF on the z/OS platform. This adapter and domain category combination uses the event hourly summarization strategy for the staging and aggregating of data. MXG software is the primary tool that is used in staging the data and in performing the hourly summarization of the staged data. A SAS IT Resource Management staging transformation uses the MXG staged output as input, and then completes the staging of the CICS transaction data. The staging job generates the XCICSUM staged table as output. The following table lists the domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.12 Domain Categories and Staged Tables Output from the Adapter Staging Job

Domain Category	Staged Table	MXG Feeder Data Sets	Adapter Setup Wizard Default Availability	Comments and Significant Changes for SAS IT Resource Management 3.2
CICS transaction activity and performance	XCICSUM	CICS	Yes - 1, 2, 3	None

This staged table is input to a number of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, and monthly time periods. Hourly and shift-based activity is included in many of these aggregations. There are no sub-hourly aggregations for this domain category because the staged data is pre-summarized to the hourly level. A filter that queries transaction ID is used to exclude all CICS management transactions from the aggregated data. Other filters can easily be created that can be used to exclude conversational transactions, and to include or exclude various operating environments (such as production, test, development, or quality assurance) from the aggregated data. These aggregations are structured by domain categories. All domain categories are further broken down into subsets of aggregations known as sub-categories. The following table summarizes these domain categories and sub-categories.

Table A7.13 Domain Categories for Staged tables (CICS Domain Category)

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
CICS transaction activity and performance	XCICSUM	CICS	CicsEnvironment	Yes - 2, 3	CICS transaction activity summarized by processing environment (that is, Production, Test, Development, Quality Assurance, and so on) as customized by the local site
			CicsApplid	Yes - 1, 2, 3	CICS transaction activity summarized by Application Region
			CicsApplid Tran	Yes - 1, 2, 3	CICS transaction activity summarized by Application Region and Transaction ID
			CicsSystem	No	CICS transaction activity summarized by Machine

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
			CicsTran	Yes - 3	CICS transaction activity summarized by Transaction ID
			CicsTranApplidUser	No	CICS transaction activity summarized by Application region, Transaction ID, and User ID
			CicsuserTran	No	CICS transaction activity summarized by User ID and Transaction ID

The following key classifications are found in this domain category:

- Processing environment

Note: The customer must customize this value by adding code or logic to the staged table Environment column that computes the valid applicable values for the local customer site.

- System (Machine)
- Application region
- Transaction ID
- User ID

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. Key metrics and specified statistics are included in special Key Metrics aggregations that are intended for longer aging periods and used for capacity planning and forecasting reporting. Ranked metrics and specified statistics are included in other aggregation tables to enable the comparison of resource usage among various classification values within the aggregated data. They are also used to filter on the top or bottom rated resource values within a set of classification values. The latter ranking filters are very useful to limit the number of reports that are created to those classifications with the greatest or the least resource utilization. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas, using the classifications previously specified:

- transaction counts
- average response time
- average and accumulated CPU time
- file access counts
- counts and percentages of selected response time criteria
- monitoring of service level objectives for transactions

Note: SAS IT Resource Management includes the new metrics that the ASUMCICX member creates. This MXG member is stored in the SAS IT Resource Management CPMISC PDS on z/OS, or the **misc** folder on UNIX, or the **sasmisc** folder on Windows. Before you use the ASUMCICX member, you must run the UTILEXCL MXG job. This job creates and implements the IMACEXCL member that is customized for your site. Documentation for ASUMCICX and UTILEXCL can be found inside each member.

SMF Adapter - DB2 Domain Category

DB2 data represents record type 101 collected by SMF on the z/OS platform. This adapter and domain category combination uses the event activity strategy for the staging and aggregating of data. MXG software is the primary tool that is used in staging the data. A SAS IT Resource Management staging transformation uses the MXG staged output as input, and then completes the staging of the DB2 accounting data. The staging job that the Adapter Setup wizard creates from the DB2 report group generates the XDBACCT staged table as output. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.14 Domain Categories and Staged Tables (DB2 Domain Category)

Domain Category	Staged Table	MSG Feeder Data Sets	Adapter Setup Wizard Default Availability	Comments and Significant Changes for SAS IT Resource Management 3.2
DB2 Batch job thread activity and performance	XDBACCT	DB2ACCT	Yes	
DB2 CICS transaction thread activity and performance	XDBACCT	DB2ACCT	Yes	
DB2 Global thread activity and performance (includes all categories of activity together, that is, batch jobs, CICS, and DDF)	XDBACCT	DB2ACCT	Yes	

This staged table is input to a number of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, and monthly time periods. Hourly and shift-based activity is included in many of these aggregations. These aggregations are structured by domain categories. All domain categories are further broken down into subsets of aggregations known as sub-categories. The following table summarizes these domain categories and sub-categories.

Table A7.15 Domain Categories for Staged Tables (DB2 Domain Category)

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
DB2 Batch Job thread activity and performance	XDBACCT	DB2 Batch	DB2BatchAttach	Yes - 3	Batch thread activity summarized by Attachment type
			DB2BatchJobAttach	Yes - 1, 2, 3	Batch thread activity summarized by Job name and Attachment type
			DB2BatchConName	Yes - 3	Batch thread activity summarized by Connection name
			DB2BatchJobbConName	Yes - 2, 3	Batch thread activity summarized by Job name and Connection name
			DB2BatchJob	Yes - 3	Batch thread activity summarized by Job name
			DB2BatchLocation	Yes - 3	Batch thread activity summarized by Location
			DB2BatchPlan	Yes - 3	Batch thread activity summarized by Plan name
			DB2BatchJobPlan	Yes - 2, 3	Batch thread activity summarized by Job name and Plan name
DB2 CICS transaction thread activity and performance	XDBACCT	DB2 Cics	DB2CicsConName	No	CICS thread activity summarized by Connection name

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
			DB2CicsCorrelate	No	CICS thread activity summarized by Correlation ID
			DB2CicsPlan	Yes - 3	CICS thread activity summarized by Plan
			DB2CicsTrans	No	CICS thread activity summarized by Transaction ID
			DB2CicsTransPlan	Yes - 2, 3	CICS thread activity summarized by Transaction ID and Plan name
DB2 DDF thread activity and performance	XDBACCT	DB2 DDG	DB2DdfAttach	Yes - 2, 3	DDF thread activity summarized by Attachment type
			DB2DdfJobAttach	Yes - 1, 2, 3	DDF thread activity summarized by Job name and Attachment type
			DB2DdfClientAppl	No	DDF thread activity summarized by Client environment and Application name
			DB2DdfNetLocation	Yes - 3	DDF thread activity summarized by Local and Remote network IDs and Location name
			DB2DdfServClasses	No	DDF thread activity summarized by Service class

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
DB2 Global thread activity and performance (includes all categories of activity, that is, batch jobs, CICS, and DDF)	XDBACCT	DB2 Global	DB2Environment	Yes - 3	DB2 thread activity summarized by processing environment (that is, Production, Test, Development, Quality Assurance, and so on) as customized by the local site
			DB2Subsystem	Yes - 1, 2, 3	DB2 thread activity summarized by subsystem
			DB2Attach	Yes - 3	DB2 thread activity summarized by Attachment type
			DB2ConType	Yes - 2, 3	DB2 thread activity summarized by Connection type
			DB2ThreadType	No	DB2 thread activity summarized by Thread type
			DB2ConName	No	DB2 thread activity summarized by Connection name
			DB2AttachConName	Yes - 1, 2, 3	DB2 thread activity summarized by Attachment type and Connection name
			DB2Correlate	No	DB2 thread activity summarized by Correlation ID

Domain Category	Staged Table Used as Input	Domain Category Name	Sub-Category Names	Adapter Setup Wizard Default Availability	Comments Specific to Sub-Category Aggregation Processing
			DB2JobAttach	Yes - 2, 3	DB2 thread activity summarized by Job name and Attachment type
			DB2Location	Yes - 3	DB2 thread activity summarized by Location name
			DB2Plan	Yes - 2, 3	DB2 thread activity summarized by Plan name
			DB2Trans	No	DB2 thread activity summarized by Transaction ID

The key classifications that are found in this domain category are:

- Processing environment

Note: The customer must customize this value by adding code or logic to the staged table Environment column in order to compute the valid applicable values for the local customer site. For optimal reporting, sites need to customize the classification column ENVIRONMENT to reflect the applicable Processing environments that are defined locally. This customization, which ideally is performed before any data collection, would typically correlate the processing environment with activity on one or more DB2 subsystems. For example, activity on DB2 subsystems DB2P and DB2A could be considered production work, so the ENVIRONMENT column could be set to 'Production' for them. Activity on DB2 subsystem DB2D could be considered development work, so the ENVIRONMENT column could be set to 'Development' for it. Activity on DB2 subsystem DB2Q could be considered quality assurance work, so the ENVIRONMENT column could be set to 'QA' or 'Quality Assurance' for it. Other criteria, along with the subsystem, might be used to establish the proper processing environment to be used in reporting. The customization of the ENVIRONMENT column is part of the staging transformation customization.

To perform the customization, perform the following steps:

1. Open the staging job that performs the staging of the DB2 accounting data (it outputs XDBACCT)
2. Right-click XDBACCT and select **Properties**.
3. Click the **Columns** tab.
4. Scroll down to the ENVIRONMENT column.

5. In the **Expression** column, replace the " " code with your replacement code that assigns the proper values for processing environment at your local site.
 6. To verify your code changes, stage some DB2 accounting data using the staging job. Then, review the output to check that the ENVIRONMENT column contains the values that you expect for your local site.
- System (Machine)
 - DB2 subsystem
 - Connection type
 - Connection name
 - Attachment type
 - Plan name
 - Location name
 - Correlation ID
 - Authorization ID
 - Job name
 - Transaction ID
 - Network ID (Local and Remote)
 - Client platform
 - Client application

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. Key metrics and specified statistics are included in special Key Metrics aggregations that are intended for longer aging periods and used for capacity planning and forecasting reporting. Ranked metrics and specified statistics are included in other aggregation tables to enable the comparison of resource usage among various classification values within the aggregated data. They are also used to filter on the top or bottom rated resource values within a set of classification values. The latter ranking filters are very useful to limit the number of reports that are created to those classifications with the greatest or the least resource utilization. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas, using the classifications previously specified:

- thread counts
- CPU, Elapsed, and I/O time in DB2
- CPU and Elapsed total time (inside and outside DB2)
- CPU and Elapsed percentages in DB2
- lock, latch, and I/O wait time in DB2
- lock escalations
- SQL operations counts

SNMP Adapter - Network Activity, Device and Interface Data, and System Data

The SNMP adapter contains several domain categories that are based on MIB II SNMP metrics. The data that is collected by default by the SNMP adapter is focused on measuring network activity through several protocols (IP, TCP, UDP, and ICMP), as well as device and interface data and system data. The SAS IT Resource Management staging transformation for SNMP is used in staging the data. The following table lists the various staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.16 Domain Categories for Staged Tables (SNMP Adapter Domain Categories for Network Activity)

Domain Category	Staged Table
Device and Interface data	SNMP_IF
System data	SNMP_System
IP protocol data	SNMP_IP
TCP protocol data	SNMP_TCP
ICMP protocol data	SNMP_ICMP
UDP protocol data	SNMP_UDP

These staged tables are input into a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, and monthly time periods. Hourly and shift-based activity is included in a number of these aggregation tables, as well as several simple aggregations.

The following key classifications are found in these domain categories:

- Device
- Interface

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- Throughput - bits per second, packets or datagrams per second, packet sizes, and packet types
- Error Rates
- Utilization and Availability by interface

- TCP Connections

VMware Adapter - Guest and Host Disk, Network, and System Domain Categories

These domain categories represent a subset of data stored in the VMware vCenter database. They focus on host and guest level activity and performance. These adapter and domain category combinations use the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.17 Domain Categories for Staged Tables (VMware vCenter Adapter)

Domain Category	Staged Table
Guest Disk	GUESTDISK
Guest Network	GUESTNETWORK
Guest System	GUESTSYSTEM
Host Disk	HOSTDISK
Host Network	HOSTNETWORK
Host System	HOSTSYSTEM

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables.

The following key classifications are found in this domain category:

- Cluster
- Guest
- Host
- DiskUnit

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- utilization of CPU and memory, CPU ready, memory ballooning, and paging for Guests
- utilization of CPU and memory, memory ballooning, and paging for Hosts

- transfer rates and counts of reads and writes for disks
- packet and data volume transfer activity for network interfaces

MS System Center Operations Manager - MS SCOM Server-Based Domain Categories

These domain categories represent a subset of data collected by a Microsoft product called System Center Operations Manager. They focus on server-level activity and performance on the Windows platform. These adapter and domain category combinations use the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.18 Domain Categories for Staged Tables (MS System Center Operations Manager for Server Domain Category)

Domain Category	Staged Table
Server Disk	SCOM_WinServerDisk
Server Network	SCOM_WinServerNwInterface
Server Paging File	SCOM_WinServerPagingFile
Server PhysicalDisk	SCOM_WinServerPhysicalDisk
Server ProcessorCpu	SCOM_WinServerProcessor
Server System	SCOM_WinServerSystem

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables.

The following key classifications are found in this domain category:

- System (Machine)
- CPU ID
- Disk ID
- Network Interface
- Paging File ID

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special KeyMetrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key

and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- utilization of CPU and memory the global server level
- utilization, response times, transfer rates, and wait and queuing for logical and physical disks

MS System Center Operations Manager Adapter - MS SCOM Exchange Domain Category

This domain category represents a subset of data that is collected by a Microsoft product called System Center Operations Manager, which focuses on Microsoft Exchange server activity and performance on the Windows platform. This adapter and domain category combination uses the server activity strategy for the staging and aggregating of data. A SAS IT Resource Management staging transformation is used in staging the data. The following table lists the various domain categories and staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.19 Domain Categories for Staged Tables (MS System Center Operations Manager Exchange Domain Category)

Domain Category	Staged Table
Exchange Disk	SCOM_ExDisk
Exchange InfoStore	SCOM_InfoStore
Exchange InfoStoreMailbox	SCOM_Mailbox
Exchange InfoStorePublic	SCOM_Public
Exchange MailServer	SCOM_SmtpServer
Exchange Network	SCOM_ ExNetworkInterface
Exchange Paging File	SCOM_ExPagingFile
Exchange PhysicalDisk	SCOM_ExPhysicalDisk

These staged tables are input to a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, monthly, and sub-hourly time periods. Hourly and shift-based activity is included in a number of these aggregation tables.

The following key classifications are found in this domain category:

- System (Machine)
- Information Store ID
- Mail Server ID

In all the domain categories that are listed in the previous table, one or more metrics are designated as key or ranked metrics. These metrics are included in special KeyMetrics

aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following areas:

- Remote Procedure Call (RPC) response times, transfer rates, and request and operation rates at the global Exchange server level
- total and active user counts at the mailboxes' information stores
- delivered, sent, and submitted messages and receive and send queue lengths for mailboxes' information stores

Web Log Adapter - All Domain Categories

The Web Log adapter contains four domain categories that are supported by the Adapter Setup wizard. These domain categories represent different categorizations of the same, limited set of metrics that are available from Web logs. The data that is collected focuses on measuring page counts and requests, the size of those requests and the response time to service those requests. The SAS IT Resource Management staging transformation for Web Log is used in staging the data. The following table lists the various staged tables that are output from the adapter staging job that is created by the Adapter Setup wizard.

Table A7.20 Domain Categories for Staged Tables (Web Log Adapter)

Domain Category	Staged Table
Browser Activity	WEBBRW
Server Activity	WEBSRV1
Visitor Organization Activity	WEBORG
Web Page Activity	WEBRSC

These staged tables are input into a series of aggregation jobs that create aggregation tables across a spectrum of daily, weekly, and monthly time periods. Hourly and shift-based activity is included in a number of these aggregation tables, but no simple aggregations are included.

The following key classifications are found in these domain categories:

- Host Name
- Site Name
- Machine
- Domain
- Browser Name
- State
- Country

- Organization Type
- URL
- Status

In all but the Browser Activity domain category, one or more metrics are designated as key or ranked metrics. These metrics are included in special Key Metrics aggregation tables that are intended for longer aging periods and used for capacity planning and forecasting reporting. They are also included in other aggregation tables as ranked metrics to enable the comparison of resource usage among a large number of servers or other classifications within the data. For specific information about the key and ranked metrics for this adapter and domain category, see [“What Are Key Metrics?” on page 417](#).

Report-ready data for this adapter focuses primarily on the following metrics:

- Bytes sent and received
- Number of pages received
- Number of requests sent
- Response time

Appendix 8

Naming Standards

About Naming Standards for SAS IT Resource Management Objects	503
Naming Standards and Location of Objects	503
Aggregations	503
Information Maps	505
Reporting	507
Staging	508

About Naming Standards for SAS IT Resource Management Objects

SAS IT Resource Management uses specific naming conventions and locations when naming and saving the various folders, jobs, transformations, filters, and tables that it supplies and creates. These conventions simplify the management and maintenance of the various objects that are interrelated. For best results, if you choose to create or edit IT data mart folders, subfolders, jobs, transformations, tables, or any other object for processing your IT resource data, then use the same naming and storage conventions that SAS IT Resource Management uses for consistency and efficiency.

Naming Standards and Location of Objects

Aggregations

The following table describes the standard naming conventions and storage locations that SAS IT Resource Management uses when creating and saving the various objects that are associated with aggregations.

Object	Naming Convention	Location	Notes
aggregation job	<domain category> Aggregation (Example: Disk Aggregation)	In the IT Data Marts tree: <IT data mart name> ⇨ <adapter type><number> ⇨ Domain Categories ⇨ <domain category> .	

Object	Naming Convention	Location	Notes
Aggregation transformation	The naming standard for aggregation transformations can vary slightly. However, the typical structure is <adapter identifier> <domain category> . (Example: NTSMF Disk)	An active Aggregation transformation is in the aggregation job. A template to create your own Aggregation transformation is in the Transformations tree in ITRM Transformations .	Transformation names cannot contain double quotation marks.
column labels (aggregation table)	<corresponding staged table column name> <statistic>	In an aggregation table.	Column names use the following guidelines: <ul style="list-style-type: none"> • Mixed case and no spaces. • Abbreviate when necessary or standardized (such as Geometric Mean to GMean or Range to Rng).
column label (ranked)	<p>In the Adapter Setup wizard: <measurement>Rank (Example: CacheReadHitPctWMeanAscRank)</p> <p>In the aggregation wizards: <measurement>Rank<number (optional)> (Examples: CacheReadHitPctWMeanDescRank, CacheReadHitPctWMeanDescRank01)</p>	In a table.	<p>In the Adapter Setup wizard:</p> <ul style="list-style-type: none"> • Mixed case. • Appended with a number. <p>In the aggregation wizards:</p> <ul style="list-style-type: none"> • Mixed case. • Appended with a number if necessary. The numbering starts with the second instance of a label with a similar name. For example, CRHTPCTRank1 has a label of CacheReadHitPctWMeanDescRank with no appended number. But CRHTPCTRank2 has a label of CacheReadHitPctWMeanDescRank01 where the appended number is 01.

Object	Naming Convention	Location	Notes
column name (ranked)	<p>In the Adapter Setup wizard: <measurement><rank><number> (Examples: DAYDATERANK01, CRHTPCTRANK01)</p> <p>In the Aggregation wizards: <measurement>Rank<order (optional)> (Examples: DAYDATEDescRank, CRHTPCTRank)</p>	In a table.	<p>In the Adapter Setup wizard:</p> <ul style="list-style-type: none"> • Uppercase. • Appended with a number. <p>In the aggregation wizards:</p> <ul style="list-style-type: none"> • Mixed case. • Not always appended with a number.
library (aggregation)	<domain category> <unique number> (Example: Disk 5455)	In the IT Data Marts tree: <IT data mart name> ⇔ <adapter type><number> ⇔ Domain Categories ⇔ <domain category> .	
libref (in an aggregation library)	AGG <same unique number as the library> (Example: AGG 5455)		
simple aggregation table	Detail<measurement category> (Example: DetailDisk)	In the IT Data Marts tree: <IT data mart name> ⇔ <adapter type><number> ⇔ Domain Categories ⇔ <domain category> .	Mixed case with no spaces.
summarized aggregation table	<time period><measurement category> (Example: MonthShiftDisk)	In the IT Data Marts tree: <IT data mart name> ⇔ <adapter type + number> ⇔ Domain Categories ⇔ <domain category> .	Mixed case with no spaces.
summarized aggregation table (key metrics)	KeyMetrics<measurement category> (Example: KeyMetricsDisk)	In the same folder as its aggregation tables and aggregation job. In the IT Data Marts tree: <IT data mart name> ⇔ <adapter type><number> ⇔ Domain Categories ⇔ <domain category> .	

Information Maps

The following table describes the standard naming conventions and storage locations that SAS IT Resource Management uses when creating and saving the various objects that are associated with information maps.

Object	Naming Convention	Location	Notes
information map	An information map has the same name as its corresponding source table. (Example: MonthShiftDisk)	In the IT Data Marts tree: <IT data mart name> ⇨ <adapter type><number> ⇨ Domain Categories ⇨ <domain category>.	
information map filter (date type)	Includes descriptive words that show the performance area first, then a verb, the type of metric, and a statistic. (Examples: MemoryUsagePctWMean, PagesFreedRateMax)	In the Information Map transformation that is part of the information map job. Information map filters are available on the Information Map Parameters tab of the Information Map transformation properties.	Mixed case with no spaces.
information map filter (statistic type)	Includes descriptive words and components that are based on the following guidelines: <ul style="list-style-type: none"> The first two characters (if the filter is based on a rank column from a joined table) show which level the column comes from (such as D_ for Day level table, W_ for Week level table, M_ for Month level table, or nothing when it is not from a joined table). The text string Top or Bot corresponds to ascending or descending order. The label name on which the column is based, when appropriate, is included. The text string DescRank or AscRank is included to show the rank order when appropriate. (Examples: BotMemoryAvailableMB, TopCpuBusyPct)	In the Information Map transformation that is part of the information map job. Information map filters are available on the Information Map Parameters tab of the Information Map transformation properties.	Mixed case with no spaces.
information map job	<domain category> Information Map (Example: Disk Information Map)	In the IT Data Marts tree: <IT data mart name> ⇨ <adapter type><number> ⇨ Domain Categories ⇨ <domain category>.	

Object	Naming Convention	Location	Notes
Information Map transformation	<time period><measurement category> (Example: MonthShiftDisk)	An active Information Map transformation is in the information map job. A template to create your own Information Map transformation is in the Transformations tree in ITRM Transformations .	<ul style="list-style-type: none"> The same name as the aggregation table that it uses. Cannot contain double quotation marks.

Reporting

The following table describes the standard naming conventions and storage locations that SAS IT Resource Management uses when creating and saving the various objects that are associated with reporting.

Object	Naming Convention	Location	Notes
domain category	Includes descriptive words for the resource that is being measured.	In the Folders tree: Products ⇒ SAS IT Resource Management ⇒ Adapters ⇒ <adapter type> ⇒ IT Template Tables .	
Gallery transformation	<domain category> <time period> (Example: Disk Monthly)	An active Gallery transformation is in the reporting job. A template to create your own Gallery transformation is in the Transformations tree in ITRM Transformations .	Transformation names cannot contain double quotation marks.
report definition	Includes descriptive words that indicate the measures that are reported, the filters that are used, and the time period that is reported. (Example: DiskBusyHighUsageHours ForTheDay)	In the Folders tree: Shared Data ⇒ SAS IT Resource Management ⇒ 3.2 IT Report Definitions ⇒ <adapter type> ⇒ <domain category or performance area> ⇒ Supplied ⇒ <time period> .	Save the report definitions that you create in the Folders tree at Shared Data ⇒ SAS IT Resource Management ⇒ IT Report Definitions ⇒ <adapter type> ⇒ <domain category or performance area> ⇒ User Defined ⇒ <time period> .
reporting job	<domain category> <time period> Reporting (Example: Disk Monthly Reporting)	In the IT Data Marts tree: <IT data mart name> ⇒ <adapter type><number> ⇒ Domain Categories ⇒ <domain category> .	

Staging

The following table describes the standard naming conventions and storage locations that SAS IT Resource Management uses when creating and saving the various objects that are associated with staging.

Object	Naming Convention	Location	Notes
column labels (staged table)	<performance area or object> <verb (if needed)> <interpretation> (Examples: MemoryUsagePct, MemoryActivePages, PagesFreedRate)	In a staged table.	Column names use the following guidelines: <ul style="list-style-type: none"> • Mixed case and no spaces. • Limit of 40 characters. • Abbreviate when necessary or standardized (such as Average to Avg, Minimum to Min, or Percent to Pct).
library (staging)	<adapter type> Staging <unique number> (Example: SAR Staging 1115)	In the IT Data Marts tree: <IT data mart name> ⇨ <adapter type><number> ⇨ Staging.	Mixed case with no spaces.
libref (in a staging library)	STG <same unique number as the library> (Example: STG 1115)		
staging job	<adapter type> Staging (Example: SAR Staging)	In the IT Data Marts tree: <IT data mart name> ⇨ <adapter type><number> ⇨ Staging.	
staged table	A staged table has the same name as the template table on which it is based.	In the IT Data Marts tree: <IT data mart name> ⇨ <adapter type><number> ⇨ Staging.	
staging transformation	<adapter type> (Example: SAR)	An active staging transformation is in the staging job. A template to create your own staging transformation for an adapter is in the Transformations tree in ITRM Transformations.	Transformation names cannot contain double quotation marks.

Appendix 9

Macros

Introduction to the Macros in SAS IT Resource Management	510
%RMDELETE	511
%RMDELETE Overview	511
%RMDELETE Syntax	511
%RMDELETE Required Arguments	511
%RMDELETE Options	512
%RMDELETE Notes	512
%RMDELETE Examples	512
%RMDUPCHK	513
%RMDUPCHK Overview	513
%RMDUPCHK Syntax	513
%RMDUPCHK Required Arguments	514
%RMDUPCHK Options	515
%RMDUPCHK Notes	516
%RMDUPCHK Example	516
%RMDUPDSN	516
%RMDUPDSN Overview	516
%RMDUPDSN Syntax	517
%RMDUPDSN Required Arguments	517
%RMDUPDSN Notes	517
%RMDUPDSN Example	517
%RMDUPINT	517
%RMDUPINT Overview	517
%RMDUPINT Syntax	517
%RMDUPINT Notes	517
%RMDUPINT Example	518
%RMDUPUPD	518
%RMDUPUPD Overview	518
%RMDUPUPD Syntax	518
%RMDUPUPD Notes	518
%RMDUPUPD Example	518
%RMMSBMCP	519
%RMMSBMCP Overview	519
%RMMSBMCP Syntax	519
%RMMSBMCP Required Arguments	519
%RMMSBMCP Options	519
%RMMSBMCP Example	520
%RMMSDTPS	521

%RMMSDTPS Overview	521
%RMMSDTPS Syntax	521
%RMMSDTPS Required Arguments	521
%RMMSDTPS Options	521
%RMMSDTPS Example	522
%RMMSSNMP	523
%RMMSSNMP Overview	523
%RMMSSNMP Syntax	523
%RMMSSNMP Required Arguments	523
%RMMSSNMP Options	523
%RMMSSNMP Example	524
%RMPOINT	524
%RMPOINT Overview	524
%RMPOINT Syntax	525
%RMPOINT Required Arguments	525
%RMPOINT Options	525
%RMPOINT Notes	525
Fast-Moving Counters	527
%RMPOINT Example	529
%RMRUNETL	529
%RMRUNETL Overview	529
%RMRUNETL Syntax	530
%RMRUNETL Required Arguments	530
%RMRUNETL Options	530
%RMRUNETL Notes	531
%RMRUNETL Examples	531

Introduction to the Macros in SAS IT Resource Management

SAS IT Resource Management provides the following macros:

- “%RMDELETE” on page 511 - for purging, erasing, or deleting an IT data mart or a table.
- “%RMDUPCHK” on page 513 - for duplicate-data checking.
- “%RMDUPDSN” on page 516 - for duplicate-data checking.
- “%RMDUPINT” on page 517 - for duplicate-data checking.
- “%RMDUPUPD” on page 518 - for duplicate-data checking.
- “%RMMSBMCP” on page 519 - for working with the BMC Perf Mgr file adapter.
- “%RMMSDTPS” on page 521 - for working with the DT Perf Sentry adapter.
- “%RMMSSNMP” on page 523 - for working with the SNMP adapter.
- %RMPDB2DM - for migrating a SAS IT Resource Management 2.7 PDB to a SAS IT Resource Management IT data mart. For more information, see *SAS IT Resource Management: Migration Documentation*. To locate the SAS IT Resource Management documentation, use the Products Index at <http://support.sas.com/documentation/index.html>.

- “%RMPROINT” on page 524 - for converting counters to rates within staging code.
- “%RMRUNETL” on page 529- for running ETL jobs.
- %RMUPGRDM - for upgrading the data from an IT data mart that is based on the data model of SAS IT Resource Management 3.1.1 or SAS IT Resource Management 2.7 to the data model of SAS IT Resource Management 3.2. For more information, see *SAS IT Resource Management: Migration Documentation*. To locate the SAS IT Resource Management documentation, use the Products Index at <http://support.sas.com/documentation/index.html>.

%RMDELETE

%RMDELETE Overview

The %RMDELETE macro enables you to purge, erase, or delete either a table or an entire IT data mart. The TYPE=, DATAMART=, and TABLEID= arguments determine the actions to perform.

%RMDELETE Syntax

```
%RMDELETE(
  DATAMART=name-of-ITdatamart | TABLEID=table-ID
  ,METAPASS=user-password
  ,METASERVER=name-of-metadata-server
  ,METAUSER=user-ID
  ,TYPE=PURGE | ERASE | DELETE
  ,WORKDIR=directory-path-for-work-files
  ,<METADATAREPOSITORY=name-of-metadata-repository>
  ,<METAPORT=port-for-metadata-server>
);
```

%RMDELETE Required Arguments

DATAMART=name-of-ITdatamart

specifies the name of the IT data mart that is to be purged, erased, or deleted. If the TABLEID argument is not specified, a value for this option is required.

TABLEID=table-ID

specifies the name of the table that is to be purged, erased, or deleted. If the DATAMART argument is not specified, a value for this option is required.

METAPASS=user-password

specifies the password that is associated with the metauser ID. A value for this parameter is required.

METASERVER=name-of-metadata-server

specifies the name of the metadata server. A value for this parameter is required.

METAUSER=user-ID

specifies the user ID that is to access the metadata server. A value for this parameter is required.

TYPE=PURGE | ERASE | DELETE

specifies the function that the macro is to perform. The function varies depending on whether an IT data mart or a table is specified. A value for this parameter is required.

- | | |
|--------|---|
| PURGE | If the DATAMART argument is specified, the PURGE function deletes the contents of all the physical tables that are associated with the specified IT data mart. If the TABLEID argument is specified, the PURGE function deletes the contents of the specified physical table. |
| ERASE | If the DATAMART argument is specified, the ERASE function deletes all the physical tables, libraries, and folders that are associated with the specified IT data mart. It also deletes the metadata for the specified IT data mart. If the TABLEID argument is specified, the ERASE function deletes the physical table and the metadata for the specified table. |
| DELETE | If the DATAMART argument is specified, the DELETE function deletes the metadata for the specified IT data mart. If the TABLEID argument is specified, the DELETE function deletes the metadata for the specified table. |

WORKDIR=directory-location

specifies a directory where %RMDELETE can write temporary external files. If this directory does not exist, it is created. For z/OS, this parameter needs to be a UNIX System Services directory path (HFS or zFS file system). A value for this parameter is required.

%RMDELETE Options**METAREPOSITORY=name-of-metadata-repository**

specifies the name of the metadata repository. The default value for this option is Foundation.

METAPORT=port-for-metadata-server

specifies the port for the metadata server. The default value for this option is 8561.

%RMDELETE Notes

On Windows, run this macro with the sasbatch.bat command in order to have the correct options set.

The %RMDELETE code must be run on the same machine that the data is located, because the data libraries are accessed with local LIBNAME statements.

If %RMDELETE is accessing an IT data mart on a metadata server on z/OS, the SAS session and the metadata server must have the same character-set encoding.

%RMDELETE Examples**Example 1: Purging All Tables in an IT Data Mart**

The following example purges all the tables in the IT data mart called “myITDataMart.”

```
%rmdelete (metaserver=myServer,
           metauser=myUserid,
```

```
metapass=myPassword,
workdir=c:\temp\workdir,
type=purge,
datamart=myITDataMart);
```

Example 2: Erasing a Table in an IT Data Mart

The following example erases the physical table with the A55DZBQ6.B70000ZC metadata ID. It also deletes the metadata for that table.

```
%rmdelete (metaserver=myServer,
           metauser=myUserid,
           metapass=myPassword,
           workdir=c:\temp\workdir,
           type=erase,
           tableID=A55DZBQ6.B70000ZC);
```

%RMDUPCHK

%RMDUPCHK Overview

The %RMDUPCHK macro checks for duplicate data and deletes it. It also builds up record counts of incoming and deleted data and datetime ranges for each system or machine. These record counts are stored in the control data set. If the control data set indicates that a gap was detected in the data, a report is generated.

The control data set is stored in the same library as the staged tables. This data set is created and managed by the %RMDUPxxx macros. (Users do not usually access this library.)

Note: For information about how to set up the %RMDUPCHK macro, see [“Implementing Duplicate-Data-Checking Macros” on page 394](#). For information about how control data sets work, see [“Control Data Sets for Duplicate-Data Checking” on page 392](#).

%RMDUPCHK Syntax

```
%RMDUPCHK(
  ENDFILE=variable-name
  ,IDVAR=variable-name
  ,SOURCE=identifier
  ,TIMESTMP=timestamp-variable-name
  <,FORCE=YES | NO>
  <,INT=interval>
  <,KEEP=number-of-weeks>
  <,TERM=YES | NO>
);
```

%RMDUPCHK Required Arguments**ENDFILE=variable-name**

specifies the name of the SAS variable that is used as the END= keyword for the SAS INFILE statement that reads the raw data.

IDVAR=variable-name

specifies the name of the SAS variable that identifies the system or machine that generated the input data.

SOURCE=identifier

specifies a unique three-character code that identifies the type of data.

TIMESTAMP=timestamp-variable-name

specifies the name of the SAS variable that contains the datetime stamp that uniquely identifies the time of the event or interval being recorded.

The SOURCE entries for the supported adapters are listed in the following table.

Table A9.1 Source Names for Each Adapter

ADAPTER	Value for the SOURCE Parameter for %RMDUPCHK
ASG TMON2CIC	TM2
ASG TMONDB2	TMD
BMC Mainview	IMF
BMC Perf Mgr	PAT
CA TMS	TMS
DT Perf Sentry	NTS
DT Perf Sentry with MXG	NTS
HP Perf Agent	(Multiple values are needed so that %RMDUPCHK can be invoked with each value.)
HP Reporter	(Multiple values are needed so that %RMDUPCHK can be invoked with each value.)
IBM DCOLLECT	DCO
IBM EREP	ERP
IBM IMS	IMS
IBM SMF	SMF
IBM TPF	TPF
IBM VMMON	VMM

ADAPTER	Value for the SOURCE Parameter for %RMDUPCHK
MS SCOM	SCO
SAP ERP	BAT, SAP, and others. (Multiple values are needed so that %RMDUPCHK can be invoked with each value.)
SAR	SAR
MS SCOM	SCO
SNMP	SNM
VMware vCenter	(Multiple values are needed so that %RMDUPCHK can be invoked with each value.)
Web Log	WWW

%RMDUPCHK Options

FORCE=YES | NO

specifies whether duplicate input data should still be processed, or whether it is a duplicate.

- FORCE=YES indicates that, if a duplicate is detected, the duplicate data should be processed.
- FORCE=NO indicates that duplicate data should not be processed. The default value for this option is **NO**.

INT=interval

represents the maximum time gap (or interval) that is to be allowed between the timestamps on any two consecutive records from the same system or machine. If the interval between the timestamp values exceeds the value of this option, then an observation with the new time range is created in the control data set. This is referred to as a gap in the data.

The value for this option must be provided in the format hh:mm, where hh represents hours and mm represents minutes. For example, to specify an interval of 14 minutes, use INT=0:14. To specify an interval of 1 hour and 29 minutes, use INT=1:29.

The default value for this option is **0:29**, or 29 minutes.

KEEP=number-of-weeks

specifies the number of weeks for which control data will be kept. Because this value represents the number of Sundays between two dates, a value of 2 (the default) results in a maximum retention period of 20 days.

The default value for this option is **2**.

TERM=YES | NO

controls whether SAS terminates if any duplicate input data is detected.

The default value of this option is **NO**.

%RMDUPCHK Notes

The Adapter Setup wizard prompts the user to specify how to handle duplicate records. Valid entries for the mode of duplicate-data checking are Inactive, Discard, Force, or Terminate.

- **Inactive:** Duplicate-data checking is not performed. No macros are executed.
- **Discard:** Duplicate-data-checking macros are executed. FORCE=NO and TERM=NO are implied.
- **Force:** Duplicate-data-checking macros are executed. FORCE=YES and TERM=NO are implied.
- **Terminate:** Duplicate-data-checking macros are executed. FORCE=NO and TERM=YES are implied.

You can change the mode of duplicate-data-checking for a table on the Properties dialog box for that table.

Note: For information about how to set up the %RMDUPCHK macro, see [“Implementing Duplicate-Data-Checking Macros” on page 394](#). For information about how control data sets work, see [“Control Data Sets for Duplicate-Data Checking” on page 392](#).

%RMDUPCHK Example

The following example provides duplicate checking for the data that is input from the NTSMF adapter:

```
%rmdupchk (
    endfile=_eof
    ,source=nts
    ,idvar=machine
    ,int=00:18
    ,keep=52
    ,timestamp=datetime
);
```

%RMDUPDSN

%RMDUPDSN Overview

The %RMDUPDSN macro generates the name of the data duplication control data set. This is a temporary SAS data set that will contain datetime ranges for the data that is being processed when duplicate-data checking is enabled. This information is also used by other duplicate-data-checking macros, such as %RMDUPCHK and %RMDUPUPD.

For supplied adapters, the %RMDUPDSN macro is automatically submitted in the staging code when duplicate-data checking is enabled. For user-written adapters, the duplicate-data checking is not automatically enabled. For information about how to enable duplicate-data checking in the user-written staging code, see [“Steps to Implement Duplicate-Data-Checking Macros” on page 395](#).

%RMDUPDSN Syntax

```
%RMDUPDSN(
SOURCE=identifier );
```

%RMDUPDSN Required Arguments

SOURCE=*identifier*

specifies a unique three-character code that identifies the type of data. It should be the same as the value that was coded for the SOURCE= parameter of the %RMDUPCHK macro.

%RMDUPDSN Notes

This macro executes only one time. It creates a global macro variable, &RMDUPDSN, that contains the name of the data duplication control data set. The global macro variable is resolved by the %RMDUPDSN macro.

%RMDUPDSN Example

This example shows the creation of a control data set, called cpnts.dsn, that is used to detect data duplication.

```
%rmdupdsn (
    source=nts
);
```

%RMDUPINT
%RMDUPINT Overview

%RMDUPINT sets up the macro definitions that are used by the other duplicate-data-checking macros. These macro definitions generate the name of the data duplication control data set.

%RMDUPINT Syntax

```
%RMDUPINT;
```

%RMDUPINT Notes

This macro requires no parameters. It contains the macro definitions and naming conventions for the control data set. The %RMDUPINT macro also defines the &RMDUPDSN macro variable that contains the data set name that is used in the %RMDUPDSN macro.

For supplied adapters, the %RMDUPINT macro is automatically submitted in the staging code when duplicate-data checking is enabled. For user-written adapters, the

duplicate-data checking is not automatically enabled. To enable duplicate-data checking, in the user-written staging code, specify the %RMDUPINT macro in front of the staging code.

For more information about how to enable duplicate-data checking in the user-written staging code, see [“Steps to Implement Duplicate-Data-Checking Macros” on page 395](#).

%RMDUPINT Example

```
%rmdupint;
```

%RMDUPUPD

%RMDUPUPD Overview

%RMDUPUPD updates the permanent control data sets with information from a temporary control data set.

%RMDUPUPD Syntax

```
%RMDUPUPD;
```

%RMDUPUPD Notes

This macro requires no parameters. It reads the temporary control data set that was built by the %RMDUPCHK macro. It splits the contents into individual control data sets, depending on the number of staging transformations that contributed to the control data set.

The individual control data sets are then merged with their corresponding permanent data sets. If, during the merging process, the time interval between any two records is greater than the allowed time interval, then both records are output. (The allowed time interval is the value that was specified by the INT= option of the %RMDUPCHK macro.)

%RMDUPUPD subsequently generates a report that informs the user of the possibility of missing data. Records that relate to data that are older than the KEEP value are deleted. This macro is executed automatically by the staging transformations of all the supplied adapters. For user-written staging transformations, this macro must be coded to execute after the staged tables have been populated.

%RMDUPUPD Example

```
%rmdupupd;
```

%RMMSBMCP

%RMMSBMCP Overview

The %RMMSBMCP macro reads the records from a BMC Perf Mgr file that was created with the DUMP HISTORY command and creates the appropriate table and column metadata for the SAS IT Resource Management BMC Perf Mgr adapter.

Note: Run this program only after you have created an IT data mart and a job that contains a BMC Perf Mgr transformation.

%RMMSBMCP Syntax

```
%RMMSDTPS(
  LOGFILE=path-and-filename-to-BMC-Perf-Mgr-log
  ,METAPASS=user-password
  ,METASERVER=name-of-metadata-server
  ,METAUSER=user-ID
  <,APPSERVER=name-of-application-server>
  <,JOBID=metadata-id-of-staging-job>
  <,CREATEMETADATA=Yes | No>
  <,FILEDEVICE=type-of-device>
  <,FILENOOFVOLUMES=number-of-volumes>
  <,FILESAPCEPRIMARY=primary-space-allocation>
  <,FILESAPCESECONDARY=secondary-space-allocation>
  <,METAPORT=metadata-server-port>
  <,METAREPOSITORY=name-of-metadata-repository>
);
```

%RMMSBMCP Required Arguments

LOGFILE=*path-and-filename-to-BMC-Perf-Mgr-log*
specifies the path to the BMC Perf Mgr log file.

METAPASS=*user-password*
specifies the password that is associated with the metauser ID.

METASERVER=*name-of-metadata-server*
specifies the name of the metadata server.

METAUSER=*user-ID*
specifies the user ID that is to access the metadata server.

%RMMSBMCP Options

APPSERVER=*name-of-application-server*
specifies the name of the application server. The default value for this option is **SASITRM**.

JOBID=*metadata-id-of-staging-job*
specifies the metadata ID of the staging job. The metadata-ID should be in the format *Axxxxxxx.Zyyyyyyy* and follow these specifications:

- *Axxxxxxx* is any set of eight alphanumeric characters preceded by the character “A.”
- *Z* is any alphabetic character.
- *yyyyyyyy* is any set of eight alphanumeric characters.

Note: This option is required if CREATEMETADATA is set to **Y**.

CREATEMETADATA=Yes | No

specifies that the macro is to create the metadata for the tables and columns in the log file.

The default value for this option is **N**, which means do not create the metadata. The value **N** only produces a report of the metadata that would have been created.

FILEDEVICE=type-of-device

specifies the type of device. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

FILENOOFVOLUMES=number-of-volumes

specifies the number of volumes. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

FILESAPCEPRIMARY=primary-space-allocation

specifies the primary space allocation. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

FILESAPCESECONDARY=secondary-space-allocation

specifies the secondary space allocation. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

METAPORT=metadata-server-port

specifies the port of the metadata server.

METAREPOSITORY=name-of-metadata-repository

specifies the name of the metadata repository that contains the IT data mart. The default value for this option is **Foundation**.

%RMMSBMCP Example

This example reads the BMC Perf Mgr log file called

c:\BMCPperfMgr\citral_Windows_format2.dat and creates the metadata for the tables and columns defined in this log file.

```
%rmmsbmcp(createMetadata=Y
           ,jobID=A5FWNTN9.BN0000RX
           ,metapass=myPassword
           ,metaserver=myServer
           ,metauser=myUserId
           ,logfile=c:\BMCPperfMgr\citral_Windows_format2.dat
           );
```

%RMMSDTPS

%RMMSDTPS Overview

The %RMMSDTPS macro reads the header records from a DT Perf Sentry log file and creates the appropriate table and column metadata for the SAS IT Resource Management DT Perf Sentry adapter.

Note: Run this program only after you have created an IT data mart and a job that contains a DT Perf Sentry transformation.

%RMMSDTPS Syntax

```
%RMMSDTPS(
  LOGFILE=path-and-filename-to-DT-Perf-Sentry-log
  ,METAPASS=user-password
  ,METASERVER=name-of-metadata-server
  ,METAUSER=user-ID
  <,APPSERVER=name-of-application-server>
  <,JOBID=metadata-id-of-staging-job>
  <,CREATEMETADATA=Yes | No>
  <,FILEDEVICE=type-of-device>
  <,FILENOOFVOLUMES=number-of-volumes>
  <,FILESAPCEPRIMARY=primary-space-allocation>
  <,FILESAPCESECONDARY=secondary-space-allocation>
  <,METAPORT=metadata-server-port>
  <,METAREPOSITORY=name-of-metadata-repository>
);
```

%RMMSDTPS Required Arguments

LOGFILE=*path-and-filename-to-DT-Perf-Sentry-log*
specifies the path to the DT Perf Sentry log file.

METAPASS=*user-password*
specifies the password that is associated with the metauser ID.

METASERVER=*name-of-metadata-server*
specifies the name of the metadata server.

METAUSER=*user-ID*
specifies the user ID that is to access the metadata server.

%RMMSDTPS Options

APPSERVER=*name-of-application-server*
specifies the name of the application server. The default value for this option is **SASITRM**.

JOBID=*metadata-id-of-staging-job*
specifies the metadata ID of the staging job. The metadata-ID should be in the format *Axxxxxxx.Zyyyyyyy* and follow these specifications:

- *Axxxxxxx* is any set of eight alphanumeric characters preceded by the character “A.”
- *Z* is any alphabetic character.
- *yyyyyyy* is any set of eight alphanumeric characters.

Note: This option is required if CREATEMETADATA is set to **Y**.

CREATEMETADATA=Yes | No

specifies that the macro is to create the metadata for the tables and columns in the log file.

The default value for this option is **N**, which means do not create the metadata. The value **N** produces only a report of the metadata that would have been created.

FILEDEVICE=type-of-device

specifies the type of device. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

FILENOOFVOLUMES=number-of-volumes

specifies the number of volumes. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

FILESAPCEPRIMARY=primary-space-allocation

specifies the primary space allocation. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

FILESAPCESECONDARY=secondary-space-allocation

specifies the secondary space allocation. This value is required only if the library and tables are going to be stored on z/OS with a traditional file system.

METAPORT=metadata-server-port

specifies the port of the metadata server.

METAREPOSITORY=name-of-metadata-repository

specifies the name of the metadata repository that contains the IT data mart. The default value for this option is **Foundation**.

%RMMSDTPS Example

This example reads the DT Perf Sentry log file called **c:\DTPerfSentry\arlington-1.smf** and creates the metadata for the tables and columns defined in this log file.

```
%rmmsdtps(createMetadata=Y
           ,jobID=A50161HR.BP0001MU
           ,metapass=myPassword
           ,metaserver=myServer
           ,metauser=myUserid
           ,logfile=c:\DTPerfSentry\arlington-1.smf
           );
```

%RMMSSNMP

%RMMSSNMP Overview

The %RMMSSNMP macro reads a management information base (MIB) definition and creates the appropriate table and column metadata and formats for the SAS IT Resource Management SNMP adapter.

Note: Run this program only after you have created an IT data mart and a job that contains an SNMP transformation.

%RMMSSNMP Syntax

```
%RMMSSNMP(
  MIBFILE=path-and-filename-of-the-MIB-definition
  ,METAPASS=user-password
  ,METASERVER=name-of-metadata-server
  ,METAUSER=user-ID
  ,<APPSERVER=name-of-application-server>
  ,<FORMATLIBPATH=path-to-library-for-formats>
  ,<JOBID=metadata-id-of-staging-job>
  ,<KEEPRECOMMENDEDONLY=Yes | No>
  ,<CREATEMETADATA=Yes | No>
  ,<METAPORT=metadata-server-port>
  ,<METAREPOSITORY=name-of-metadata-repository>
);
```

%RMMSSNMP Required Arguments

MIBFILE=*path-and-filename-of-the-MIB-definition*
specifies the path to the management information base (MIB) definition file.

METAPASS=*user-password*
specifies the password that is associated with the metauser ID.

METASERVER=*name-of-metadata-server*
specifies the name of the metadata server.

METAUSER=*user-ID*
specifies the user ID that is to access the metadata server.

%RMMSSNMP Options

APPSERVER=*name-of-application-server*
specifies the name of the application server.
The default value for this option is **SASITRM**.

CREATEMETADATA=Yes | No
specifies that the macro is to create the metadata for the MIB tables.
The default value for this option is **N**, which means do not create the metadata. The value **N** only produces a report of the metadata that would have been created.

FORMATLIBPATH=*path-to-library-for-formats*

specifies the path to the library where the formats for this macro are to be stored.

The default value for this option is the WORK library.

JOBID=*metadata-id-of-staging-job*

specifies the metadata ID of the staging job. The metadata-ID should be in the format *Axxxxxxx.Zyyyyyyy* and follow these specifications:

- *Axxxxxxx* is any set of eight alphanumeric characters preceded by the character “A.”
- *Z* is any alphabetic character.
- *yyyyyyy* is any set of eight alphanumeric characters.

Note: This option is required if CREATEMETADATA is set to **Y**.

KEEPRECOMMENDEDONLY=Yes | No

specifies that the macro should keep only the recommended tables and columns.

The default value for this option is **N**, which means that all tables and columns are to be kept, including the tables and columns that are deprecated or obsolete, and columns whose Max-Access attribute has a value of **not-accessible**.

METAPORT=*metadata-server-port*

specifies the port of the metadata server.

The default value for this option is **8561**.

METAREPOSITORY=*name-of-metadata-repository*

specifies the name of the metadata repository that contains the IT data mart.

The default value for this option is **Foundation**.

%RMMSSNMP Example

This example reads the BGP-4-MIB definition and creates the metadata for the recommended tables and columns that are in the MIB definition.

```
%rmmssnmp(createMetadata=Y
           ,formatLibPath=c:\mibFormats
           ,jobID=A50161HR.BP0001MU
           ,keepRecommendedOnly=Y
           ,metapass=myPassword
           ,metaserver=myServer
           ,metauser=myUserid
           ,mibfile=c:\mibs\BGP4-MIB.my
           );
```

%RMPROINT**%RMPROINT Overview**

The %RMPROINT macro creates an output SAS view that converts the values of the variables of data type C2RATE or D2RATE from counters or deltas into rates. This macro processes either counters or deltas in a single invocation (but not both). If

DURATION is missing, then the macro also sets the value of DURATION to be the difference between successive values of the variable DATETIME within one BY group.

%RMPROINT Syntax

```
%RMPROINT(
  ADAPTER=adapter-name
  ,BYGROUP=list-of-BY-vars
  ,IN=input-SAS-table
  ,OUT=output-SAS-table
  ,VARS=list-of-vars
  <,CONVERT=C2RATE>
  <,_RBDF=relative-bad-delta-factor>
);
```

%RMPROINT Required Arguments

ADAPTER=*adapter-name*
specifies the name of the adapter to which the table belongs.

BYGROUP=*bygroup*
specifies the space-delimited list of BY variables that govern how this data is sorted.

IN=*infile*
specifies the input SAS table.

OUT=*outfile*
specifies the output SAS table. (This table can be the same as the input table.)

VAR**S**=*variable-name*
specifies the space-delimited list of variables that require conversion.

%RMPROINT Options

CONVERT=*conversion-type*
specifies the type of conversion to be performed. The following values are valid for *conversion-type*:

- C2RATE
The data is converted from counter to rate. The default value for this option is C2RATE.
- D2RATE
The data is converted from the delta value to rate.

_RBDF=*relative-bad-delta-factor*
specifies the name of the SAS macro variable that is created that contains the Relative Bad Delta Factor as a percentage. The default value of this option is RBDFTMP.

%RMPROINT Notes

The %RMPROINT macro converts metrics that are considered counters or deltas to rates. A counter is a metric whose value is constantly increasing. A delta is a metric whose value represents the difference between the value of a counter at the start and end

of an interval of time. The rate is determined by taking the delta metric (or by calculating the delta from the counter) and dividing it by the duration of the interval. The following table lists the counters that are converted to rates for the NTSMF Adapter.

Table A9.2 *List of Counters That Are Converted to Rates for NTSMF Adapters*

NTSMCNT Table	NTTCP Table
ALMSGKB	CNCACTV
EXTMSGGS	CNCFLRS
EXTRMKB	CNCTPSV
INTMSGGS	CNCTRST
INTRMKB	
TTLMSGGS	

The following table lists the counters that are converted to rates for the SAR Adapter.

Table A9.3 *List of Counters That Are Converted to Rates for SAR Adapters*

XYTCPS TABLE			
ICIMDEU	ICOMDEU	IMFQGOK	IMRMSOK
ICIMECQ	ICOMECQ	IMFGCRE	IMRMSRE
ICIMECR	ICOMECR	IMFGFLD	IMRSMAC
ICICERR	ICOMERR	IMIADRE	IMRSMFU
ICIMIPP	ICOMIPP	IMIDELV	IMRSMTO
ICIMISQ	ICOMISQ	IMIDISC	IMRTDIS
ICIMITE	ICOMITE	IMIFWDD	IMSMFLD
ICIMMKQ	ICOMMKQ	IMIHDRE	INSEGS
ICIMMKR	ICOMMKR	IMIRECV	OUTSEGS
ICIMRED	ICOMRED	IMIUNPR	USINDGR
ICIMSGS	ICOMSGS	IMODISC	USINERR
ICIMTIQ	ICOMTIQ	IMONORT	USNOPRT
ICIMTIR	ICOMTIR	IMOREQS	USOTDGR

For example, if you have recorded odometer mileage and the start time of each mileage reading, then the counter is mileage. When this macro is executed, the duration is determined by calculating the difference between the start times of one observation and the preceding observation. Then, the number of miles per second can be calculated for each observation. To calculate the number of miles per second, find the difference between the odometer mileage of one observation and the mileage of the preceding observation, and then divide the result by the duration in seconds.

A counter continues increasing across all intervals until it reaches its maximum value and then, typically, it resets to zero and starts again. The counter maximums are 65,536 (for a 16-bit counter), 4,294,967,296 (for a 32-bit counter), and 18,446,744,073,709,551,616 (for a 64-bit counter). Using the previous and current value of the counter and knowing the maximum limit for that particular counter, it is possible to calculate the rate for that interval, based on these assumptions:

- If the previous value is smaller than the current value, it is assumed that the counter has not reset.
- If the previous value is larger than the current value, it is assumed that the counter has reset one time.

Note: For information about “Length and Precision of Variables” when using any 64-bit counter, see the SAS Companion documentation for your host.

Fast-Moving Counters

A problem that exists, especially for 16-bit counters, is that fast-moving counters can reset more than one time during an interval. One way to resolve this problem is to sample at a shorter interval so that there is not enough time during the interval for the counters to reset more than once.

SAS IT Resource Management also provides another way to resolve this problem. Although SAS IT Resource Management converts counters to rates automatically, there are three macro variables that you can use to affect how the conversion is done. By default, these macro variables are set as follows:

- Macro variable CP16PCT is for use with 16-bit counters. By default, it is set as follows:

```
%let cp16pct = 0.95;
```

As a result of this setting, a threshold is calculated for 16-bit counters, and the threshold has the value of: $0.95 * 65536 = 62259.2$.

- Macro variable CP32PCT is for use with 32-bit counters. By default, it is set as follows:

```
%let cp32pct = 0.97;
```

As a result of this setting, a threshold is calculated for 32-bit counters, and the threshold has the value of $0.97 * 4292967296 = 4164178277$.

- Macro variable CP64PCT is for use with 64-bit counters. By default, it is set as follows:

```
%let cp64pct = 0.97;
```

As a result of this setting, a threshold is calculated for 64-bit counters, and the threshold has the value of $0.97 * 18446744073709551616 = 17893341751498265067$.

When a rate is being determined and the current value is less than the previous value, then one of the following situations might have occurred:

- If the previous value is greater than or equal to the threshold, the rate is calculated based on the assumption that one reset occurred during the interval. That is, the previous value was so close to the counter's maximum that it is possible that it was reset once during the interval.

In this case, a "Corrected" message (see the following example) is written to the SAS log.

- If the previous value is less than the threshold, the rate is not calculated, and, instead, the rate is set to missing. That is, the previous value is so far from the counter's maximum that the counter must have been moving extremely fast to reset during the interval. And if the counter is moving that fast, it is possible that it was reset more than one time during that interval.

In this case, a "Set to missing" message (see the following example) is written to the SAS log.

Example

The following example illustrates the situation for a 16-bit counter, assuming 5-minute intervals:

Table A9.4 Table of Counters and Calculated Rates

Counter Value	Calculated Rate
20000	.
40000	66.67
62260	72.20
300	11.92
4000	132.33
62259	74.19
2000	.

In the preceding table, the calculated rates for the first three counter values are computed according to the following rules:

- The first observation's rate is missing, because there is no previous value with which to calculate the difference.
- The first time the counter resets, the value preceding the reset (62260) meets (or exceeds) the default threshold of 95% for 16-bit counters (62259.2). The rate is therefore calculated based on one reset, and the following message is written to the SAS log:

```
(CLEANUP)Obs4 16 bit Overflow Start 01JAN01:00:10:00
SVAL=62260 Corrected End 01JAN2001:00:15:00.00 SVAL=300
```

- The second time the counter resets, the value preceding the reset (62259) does not meet (or exceed) the default threshold of 95% for 16-bit counters (62259.2). The rate is therefore set to missing, and the following message is written to the SAS log:

```
(CLEANUP)Obs 7 Inconsistent Start 01JAN01:00:25:00
SVAL=62259 Set to missing End 01JAN2001:00:30:00.00
SVAL=2000
```

If a large percentage of your data is producing the second type of message, the work that is measured by the counters might have increased so much that the counters are resetting more than one time within an interval. Take one or both of the following actions as needed:

- Reduce the interval at which you sample, although this will increase the volume of data that you collect.
- Lower the thresholds. For example, you could insert the following SAS statement in the process-and-reduce job before the call to the process macro:

```
%let cp16pct=0.94;
```

Note: As you lower the thresholds, more rates are calculated, but the rates might be incorrect because of multiple resets due to fast-moving counters.

%RMPOINT Example

If you are working with a supported adapter, this macro is submitted automatically when necessary. If you are working with user-written staging code, you must submit the macro yourself if it is needed. To do so, execute it as a separate step in a job after the data has been staged.

```
DATA
....

staging code goes here
.....
run;
%rmprint(...);
```

%RMRUNETL

%RMRUNETL Overview

The %RMRUNETL macro runs SAS IT Resource Management ETL jobs. The staging job is run first, and then the aggregation jobs are run in parallel using Multi-Processing (MP) Connect, followed by the information map jobs that are run in parallel, and finally the report jobs that are run in parallel. If the staging job fails, the aggregation, information map, and report jobs are not run. After each job completes, the SAS log is examined by the program. The completion status of the job (Success, Warning, or Error) is sent by e-mail to the specified recipient. If the log displays warnings or errors, the e-mail includes the text of the error or warning along with the corresponding line number in the log file.

%RMRUNETL Syntax

```

%RMRUNETL(
  <ADAPTER=adapter-name>
  ,<AGGJOB1=path-to-aggregation-job1>,...,<AGGJOB20=path-to-aggregation-job20>
  ,<AGGLOG1=path-to-aggregation-log1>,...,<AGGLOG20=path-to-aggregation-log20>
  ,<EMAILFROM=sending-email-address>
  ,<EMAILTO=address-where-email-is-to-be-sent>
  ,<INFOMAPJOB1=path-to-information-map-job1>,...,<INFOMAPJOB20=path-to-information-map-job20>
  ,<INFOMAPLOG1=path-to-information-map-log1>,...,<INFOMAPLOG20=path-to-information-map-log20>
  ,<REPORTJOB1=path-to-report-job1>,...,<REPORTJOB20=path-to-report-job20>
  ,<REPORTLOG1=path-to-report-log1>,...,<REPORTLOG20=location-of-report-log20>
  ,<RUNDATE=date-that-this-macro-is-run>
  ,<STAGINGJOB=path-to-staging-job>
  ,<STAGINGLOG=path-to-staging-log>
);

```

%RMRUNETL Required Arguments

There are no required parameters for this macro. All parameters are optional. However, you must specify at least one job in order to run this macro.

%RMRUNETL Options**ADAPTER=*adapter-name***

This value is used in the subject of the e-mail messages that are sent out and in the name of the audit table. This option can be left blank.

AGGJOB1=*path-to-aggregation-job1*...AGGJOB20=*path-to-aggregation-job20*

specifies the paths to a maximum of 20 aggregation jobs. These jobs will be run in parallel.

Note: The path must include the name of the file that contains the job.

AGGLOG1=*path-to-aggregation-log1*...AGGLOG20=*path-to-aggregation-log20*

specifies the paths to a maximum of 20 aggregation logs.

Note: The path must include the name of the file that contains the log.

EMAILFROM=*sending-email-address*

specifies the e-mail address from which the e-mail is to be sent. If blank, this value is set to the address that is specified in EMAILTO.

EMAILTO=*address-where-email-is-to-be-sent*

specifies the e-mail address to which the status e-mail is to be sent. If this value is blank, a warning is written in the log that e-mail is not going to be sent.

INFOMAPJOB1=*path-to-information-map-job1*...INFOMAPJOB20=*path-to-information-map-job20*

specifies the paths to a maximum of 20 information map jobs. These jobs will be run in parallel.

Note: The path must include the name of the file that contains the job.

INFOMAPLOG1=*path-to-information-map-log1*...INFOMAPLOG20=*path-to-information-map-log20*

specifies the paths to a maximum of 20 information map logs.

Note: The path must include the name of the file that contains the log.

REPORTJOB1=*path-to-report-job1*...**REPORTJOB20=***path-to-report-job20*
specifies the paths to a maximum of 20 report jobs. These jobs will be run in parallel.

Note: The path must include the name of the file that contains the job.

REPORTLOG1=*path-to-report-log1*...**REPORTLOG20=***path-to-report-log20*
specifies the paths to a maximum of 20 report logs.

Note: The path must include the name of the file that contains the log.

RUNDATE=*date-that-this-macro-is-run*
specifies the date that is used in e-mail messages and in log filenames. If this value is blank, it is set to today's date.

STAGINGJOB=*path-to-staging-job*
specifies the path to the staging job that you want to run.

Note: The path must include the name of the file that contains the job.

STAGINGLOG=*path-to-staging-log*
specifies the path to the staging log that is produced by the staging job. If the STAGINGJOB option is specified, then the STAGINGLOG must also be specified.

Note: The path must include the name of the file that contains the job.

%RMRUNETL Notes

This macro can be run in batch mode, or interactively from the SAS Display Manager. It can run Windows, UNIX, or z/OS jobs that are stored in a PDS.

If an option for a job is specified, the corresponding log option must also be specified. For example, if there is a path specified for the AggJob2 option, the AggLog2 option cannot be left blank.

%RMRUNETL Examples

Example 1: Running a Staging Job, Two Aggregation Jobs, and a Report Job

This example shows the code that runs a staging job, two aggregation jobs, and a report job. E-mail is to be sent to and from the same e-mail address. The RUNDATE option defaults to today's date.

```
%rmrunetl(adapter=SNMP
           ,stagingJob=C:\jobs\staging.sas
           ,stagingLog= C:\jobs\logs\staging.log
           ,aggJob1=C:\jobs\_1_aggJob.sas
           ,aggJob2=C:\jobs\_2_aggJob.sas
           ,aggLog1= C:\jobs\logs\_1_aggJob.log
           ,aggLog2= C:\jobs\logs\_2_aggJob.log
           ,reportJob1= C:\jobs\_1_reportJob.sas
           ,reportLog1= C:\jobs\_1_reportJob.log
           ,emailto=email.address@company.com
           );
```

Note: Information Map jobs do not have to be run unless a change was made to the maps that are associated with those jobs.

Example 2: Running an MS SCOM Job on Windows

MS_SCOM_Overall.bat: This code shows the Windows batch job that runs the MS_SCOM_Overall.sas program.

```
echo off

set ADAPTER=MSSCOM
set DATETIME=%date:~10,4%%date:~4,2%%date:~7,2%_%time:~0,2%%time:~3,2%
set BASE=C:\Data\Jobs\%ADAPTER%
set LOGDIR=%BASE%\logs
set JOBDIR=C:\SAS\ITRM32\Lev1\SASITRM\SASEnvironment\SASCode\Jobs
set SRC=MS_SCOM_Overall
set SAS=C:\SAS\ITRM32\Lev1\SASITRM\sas.bat

%SAS% -sysin %BASE%\%SRC%.sas -log %LOGDIR%\%SRC%_%DATETIME%.log
```

MS_SCOM_Overall.sas: This code shows the SAS program that runs all the daily jobs (staging, aggregation, and reporting) for an MS SCOM data mart.

```
%let adapter = %sysget(ADAPTER);
%let datetime = %sysget(DATETIME);
%let JOBhome = %sysget(JOBDIR);
%let LOGhome = %sysget(LOGDIR);

%rmrunetl(adapter = &ADAPTER
, stagingJob = &JOBhome\MS_SCOM_Staging.sas
, stagingLog = &LOGhome\ MS_SCOM_Staging_&datetime..log
, aggJob1 = &JOBhome\InfoStore_Aggregation.sas
, aggJob2 = &JOBhome\InfoStoreMailbox_Aggregation.sas
, aggJob3 = &JOBhome\Disk_Aggregation.sas
, aggJob4 = &LOGhome\System_Aggregation.sas
, aggLog1 = &LOGhome\InfoStore_Aggregation_&datetime..log
, aggLog2 = &LOGhome\InfoStoreMailbox_Aggregation_&datetime..log
, aggLog3 = &LOGhome\Disk_Aggregation_&datetime..log
, aggLog4 = &LOGhome\System_Aggregation_&datetime..log
, reportJob1 = &JOBhome\InfoStore_Daily_Reporting.sas
, reportJob2 = &JOBhome\InfoStoreMailbox_Daily_Reporting.sas
, reportJob3 = &JOBhome\Disk_Daily_Reporting.sas
, reportJob4 = &JOBhome\System_Daily_Reporting.sas
, reportLog1 = &LOGhome\InfoStore_Daily_Reporting_&datetime..log
, reportLog2 = &LOGhome\InfoStoreMailbox_Daily_Reporting_&datetime..log
, reportLog3 = &LOGhome\Disk_Daily_Reporting_&datetime..log
, reportLog4 = &LOGhome\System_Daily_Reporting_&datetime..log
, emailto = email.address@company.com
);
```

Example 3:

HP_Perf_Agent_Overall: This code shows the UNIX script (Bash, Bourne and related shells) that runs the HP_Perf_Agent_Overall.sas program.

```
export ADAPTER=HPPerfAgent
export BASE=/var/opt/sas/prog/prod/$ADAPTER
export LOGDIR=$BASE/logs
export JOBDIR=/opt/sas/ITRM/itrm32srv/Lev1/SASITRM/SASEnvironment/SASCode/Jobs
export SRC=HP_Perf_Agent_Overall
export DATETIME=$(date +"%Y%m%d_%H%M") ADAPTER=HPPerfAgent
```



```
SAS=/opt/sas/ITRM/SASHome/SASFoundation/9.3/sas
```

```
$SAS -sysin $BASE/$SRC.sas -log $LOGDIR/"$SRC_"$DATETIME".log -noterminal -rsasuser &
```

HP_Perf_Agent_Overall.sas: This code shows the SAS program that runs all the daily jobs (staging, aggregation, and reporting) for an HP Perf Agent data mart.

```
%let adapter = %sysget(ADAPTER);
%let datetime = %sysget(DATETIME);
%let JOBhome = %sysget(JOBDIR);
%let LOGhome = %sysget(LOGDIR);

%rmrunetl(adapter = &adapter
, stagingJob = &JOBhome/HP_Perf_Agent_Staging.sas
, stagingLog = &LOGhome/HP_Perf_Agent_Staging_&datetime..log
, aggJob1 = &JOBhome/DiskUnix_Aggregation.sas
, aggJob2 = &JOBhome/SystemUnix_Aggregation.sas
, aggLog1 = &LOGhome/DiskUnix_Aggregation_&datetime..log
, aggLog2 = &LOGhome/SystemUnix_Aggregation_&datetime..log
, reportJob1 = &JOBhome/DiskUnix_Daily_Reporting.sas
, reportJob2 = &JOBhome/SystemUnix_Daily_Reporting.sas
, reportLog1 = &LOGhome/DiskUnix_Daily_Reporting_&datetime..log
, reportLog2 = &LOGhome/SystemUnix_Daily_Reporting_&datetime..log
, emailto = email.address@company.com
);
```


Appendix 10

Best Practices and Troubleshooting Tips

How to Back Up SAS IT Resource Management	536
Introduction to the Backup Process	536
Scenario: Recovery of Metadata, Data, and Reports After a Job Fails	536
Backing Up Metadata	537
Backing Up Data	537
Backing Up Reports	538
Backing Up Other Reporting-Related Content	538
Debugging Tips	539
Problems with the SAS IT Resource Management Client	539
Server Logging Files	539
Turning On Standard Execution-time Debugging (ETLS_DEBUG)	539
Turning On Additional Execution-time Debugging (LOGGERLEVEL)	539
Investigating Problems Using Java from SAS Foundation	540
Tips for Working with SAS IT Resource Management Client	541
Tips for Working with the IT Data Mart	541
Managing IT Data Marts	541
Archiving Your Data	541
Storing Data Tables on z/OS Systems	542
Deleting IT Data Marts	542
Tips for Working with Aggregations	542
Tips for Working with Information Maps	543
Using SAS Information Map Studio	543
Tips for Installing and Configuring SAS IT Resource Management	544
Installing and Configuring SAS IT Resource Management	544
Handling Holidays	544
Overview of Handling Holidays	544
Example of Handling Holidays	547
Performance Issues	549

How to Back Up SAS IT Resource Management

Introduction to the Backup Process

The backup process takes a snapshot of the current state of the SAS IT Resource Management system. It is useful for record-keeping purposes to take backups before you make major changes to the tables and transformations in your system. Another good practice is to run the backup process nightly either before or after the system is being updated with new data.

Backing up your SAS IT Resource Management system consists of backing up four separate areas of content:

- Metadata
- Data
- Reports
- Other miscellaneous content

TIP For best results, establish a formal, regularly scheduled backup process. It is important to back up all the above items at (approximately) the same point in time so that related information is synchronized. This practice is crucial if a restore becomes necessary.

Make sure that a backup process is put into place immediately after SAS IT Resource Management is initially deployed, and that backups are taken before major changes are made to the system. This practice will help the data administrator to recover your system in the event of an error.

For information about backing up your system, see *SAS(R) 9.3 Intelligence Platform: System Administration Guide*, at <http://support.sas.com/documentation/index.html>.

Scenario: Recovery of Metadata, Data, and Reports After a Job Fails

In the following scenario, someone updates the metadata server and inadvertently corrupts it. The error is not discovered until the next day's processing of the daily job fails. The regular backup schedule that is shown here enables the data administrator to restore the metadata server. It also shows how a good backup schedule can be used successfully, avoiding the reloading of multiple days of data.

- Day 1
 - At 2:00 a.m., the daily job is started.
 - At 4:00 a.m. (or as soon as the daily job finishes), the backup job runs, creating generation N of the metadata, data, and reports.
 - Later in day 1, the metadata server is corrupted. The problem is not discovered at this time.
- Day 2
 - At 2:00 a.m., the daily job is started.
 - At 4:00 a.m. (or as soon as the daily job finishes), the backup job runs, creating generation N + 1 of the metadata, data, and reports.

- Later in day 2, the error in the metadata server is discovered. The system administrator can restore the metadata, data, and reports from the backup that was made on day 1 (generation N), and rerun the daily job.

Note: The same schedule of regular backups for the metadata server that is described in this scenario should be used for backing up data, reports, and other report-related objects.

Backing Up Metadata

Metadata for SAS IT Resource Management is stored in the SAS Metadata Server. The following is a partial list of some operations that result in metadata being updated:

- An IT data mart is created or modified.
- The Adapter Setup wizard is executed.
- Staged tables or aggregations tables are created or modified.
- Jobs are created or modified.
- Reports are published with the ITRM Gallery task in SAS Enterprise Guide.
- Information maps are created or modified.

The data administrator should back up the metadata either before or after these or similar activities.

Note: Metadata is not updated during the execution of the daily ETL or report jobs.

The SAS 9.3 Metadata Server includes a new server-based facility that performs metadata server backups automatically on a scheduled basis. This facility can also be used to perform ad hoc backups and roll-forward recovery. The facility replaces the %OMABAKUP macro, which was available in SAS 9.1 and 9.2, and the Backup and Restore Wizard in SAS Management Console, which was available in SAS 9.2. For complete instructions for backing up the SAS Metadata Server, see Part 4, “Backing Up and Restoring Your SAS Content,” in the *SAS 9.3 Intelligence Platform: System Administration Guide*, at <http://support.sas.com/documentation/index.html>.

TIP If you stop the SAS Metadata Server to do a backup, any servers in your system that access the Metadata Server might be adversely affected.

Backing Up Data

SAS IT Resource Management stores data in files that reside in the customer’s physical file system. These files are pointed to by metadata objects in the IT data mart. Typically, all physical files for a single IT data mart reside underneath the root path of an IT data mart. You can use the Base SAS DATASETS or COPY procedures, operating systems commands, or third-party tools to copy all of the files and subdirectories beneath this root path to a backup location. This backup procedure should be done for all IT data marts.

For best results, you should back up the entire contents of the IT data mart, instead of selecting specific files or directories. In other words, make sure that your backup procedure includes the aggregation tables, staged tables, and the ADMIN and SPIN libraries.

Backing Up Reports

SAS IT Resource Management uses SAS Content Server, a WebDAV server, to store report output. Report output is created by the report jobs and is displayed in Gallery Manager. In addition, other SAS products that are included with SAS IT Resource Management, such as SAS Web Report Studio and SAS Information Delivery Portal, use SAS Content Server for storage of their reports and supporting files.

To back up SAS Content Server, perform the following steps:

1. Stop either the Web application server or the SAS Content Server application.
 - If you are using a JBoss Application Server, then you must stop that Web application server.
 - If you are using an IBM WebSphere Application Server Network Deployment or an Oracle WebLogic Server, then you must stop the SAS Content Server application.

Note: If you stop the SAS Content Server to do a backup, any servers in your system that access the SAS Content Server might be adversely affected.

2. Use operating systems commands or third-party tools to copy all the files and subdirectories from the middle tier platform to a backup location.
 - If your middle tier platform is on Windows, back up the files and subdirectories that are in this location: `<SAS-configuration-directory>\Lev1\AppData\SASContentServer\Repository`.
 - If your middle tier platform is on UNIX, back up the files and subdirectories that are in this location: `<SAS-configuration-directory>/Lev1/AppData/SASContentServer/Repository`.
 - If your middle tier platform is on z/OS, back up the files and subdirectories that are in this location: `<SAS-configuration-directory>/Lev1/AppData/SASContentServer/Repository`.
3. Start your Web application server or SAS Content Server application.

Backing Up Other Reporting-Related Content

If updates are made to any copies of SAS Enterprise Guide projects that are used for reporting in SAS IT Resource Management, then those user-owned copied projects should be backed up by using operating system commands or third-party tools. If you installed the supplied SAS Enterprise Guide projects, they can be accessed by navigating to the location where you installed SAS IT Resource Management. The projects are located at `SASITResourceManagementSASEnterpriseGuideComponents\3.21`. For example, if you installed SAS IT Resource Management on a 64-bit machine, the projects might be located at `C:\Program Files\SASHome\x86\SASITResourceManagementSASEnterpriseGuideComponents\3.21`.

Note: Making changes to supplied SAS Enterprise Guide projects is strongly discouraged. If you must change a project, copy the project and make your changes to that copied version of the project.

Users of Gallery Manager can create galleries and gallery folders. You might want to back up those items so that they do not have to be manually recreated in the unlikely event of a system failure. If so, you should use operating systems commands or third-party tools to copy all of the files on the middle tier platform:

- If your middle tier platform is on Windows, back up the files that are in this location:
`<SAS-configuration-directory>\Lev1\AppData\ITRMGalleryManager.`
- If your middle tier platform is on UNIX, back up the files that are in this location:
`<SAS-configuration-directory>/Lev1/AppData/ITRMGalleryManager.`
- If your middle tier platform is on z/OS, back up the files that are in this location:
`<SAS-configuration-directory>/Lev1/AppData/ITRMGalleryManager.`

Debugging Tips

Problems with the SAS IT Resource Management Client

The IT Resource Management Client automatically keeps log files for certain operations such as the following:

- creating an IT data mart
- running the Adapter Setup wizard
- adding staged tables

These log files can be useful to you or to Technical Support personnel, when investigating a problem.

On the client, these log files are located in a directory that is determined based on the APPDATA environment variable.

- On Windows XP, this location is typically `C:\Documents and Settings\userid\Application Data\SASHome\SASITResourceManagement\3.21.`
- On Windows 7, this location is typically `C:\Users\userid\AppData\Roaming\SASHome\SASITResourceManagement\3.21.`

Server Logging Files

All servers, such as the SAS Metadata Server, the object spawner, and the application and workspace servers, can be configured to have logs. These logs can be useful in diagnosing problems. For information about configuring logging for these servers, see “Administering Logging for SAS Servers” in the *SAS Intelligence Platform: System Administration Guide*.

Turning On Standard Execution-time Debugging (ETLS_DEBUG)

To turn on execution-time debugging in the generated code for SAS Data Integration Studio and SAS IT Resource Management transformations, add the following line to the generated code, before the transformation for which you would like additional debugging information: `%LET ETLS_DEBUG=1;`

Turning On Additional Execution-time Debugging (LOGGERLEVEL)

To turn on additional execution-time debugging information in the generated code for SAS IT Resource Management transformations, add or edit the following line to the generated code before the transformation for which you would like additional debugging information: `%LET LOGGERLEVEL= [level] ;`

Legal values for [level] include: FATAL, ERROR, WARN, INFO, DEBUG, or TRACE. Users who want additional debugging information should set the level to **DEBUG: %LET LOGLEVEL=DEBUG;**

Investigating Problems Using Java from SAS Foundation

Some features that use Java from SAS Foundation. These features include PROC INFOMAPS, the Information Map LIBNAME engine, and the migration of IT Resource Management 2.7 PDBs to IT data marts supported in SAS IT Resource Management 3.2 and later. For problems in running Java from SAS, use the following techniques to investigate whether SAS is configured correctly to use Java:

- Make sure that the Java JVM can be started, and that it is the correct version of the JVM. To do so, submit this SAS statement:

```
proc javainfo all;
run;
```

The results of this code will include JVM details and properties. This statement can be added to the SAS code of a deployed job, or submitted directly from the **Tools** ⇒ **Source Editor** menu item in SAS Data Integration Studio.

- Check the **JREOPTIONS** settings to make sure that the Java options are correct, particularly in these elements:
 - Memory options such as **-Xmx** are specified according to the installation and configuration documents.
 - To see the current **JREOPTIONS**, submit this SAS statement:

```
proc options
option=jreoptions;
run;
```

This statement can be added to the SAS code of a deployed job, or submitted directly from the **Tools** ⇒ **Source Editor** menu item in SAS Data Integration Studio.

Note: When running jobs in SAS Foundation, make sure that SAS is invoked using the appropriate configuration files because these files have been written specifically for use by SAS IT Resource Management.

Troubleshooting Error Messages in the SAS Log

Error message is displayed in the SAS log: “ERROR: Failed to insert filter” or “ERROR: Failed to insert data source.”

Make sure that the user who is running the job has the proper authorizations. To do so, open the SAS Management Console and perform the following steps:

1. Set the repository to **Foundation**.
2. In the **User Manager** component, select the user profile of the person who is running the job.
3. On the **Logins** tab, locate the entry for the authentication domain of the workspace server that was in use at the time of the error. (Each workspace server has an authentication domain.) The user ID and the password should be specified. The password cannot be left blank.

Note: Whenever the server is reset, the password might not have been saved. (A password field of asterisks does not necessarily contain your password.) In this case, you should manually update your password.

Tips for Working with SAS IT Resource Management Client

On any **Properties** dialog box, after you make a change to a field, use the **Tab** key to move out of that field. Then click **Apply** to make sure your changes are reflected in the metadata.

Tips for Working with the IT Data Mart

Managing IT Data Marts

Here are some tips that can help you manage IT data marts effectively:

- Do not create IT data marts with duplicate names. If IT data marts do not have unique names, generated information maps might not be placed in the correct job.
- Store all data and information maps in the IT data mart.
- You can create IT data marts in different folders. The choice of a parent folder should reflect your intended use. For example, if you intend to share access to a new IT data mart, the **/Shared Data/SAS IT Resource Management/IT Data Marts** location is a logical place to create it. The **/Products** and **/System** folders are intended for product configuration files. Users are not typically permitted to write to these folders. Therefore, do not create IT data marts in those folders.
- The organization of your site's data into IT data marts depends on your site's needs. You do not have to put data from multiple adapters into the same IT data mart, and you are not prevented from doing so. If you decide to use multiple IT data marts, you can still combine data from the various IT data marts using information maps or other similar techniques.

Archiving Your Data

SAS IT Resource Management does not provide an automated archiving feature. However, archiving is easy to accomplish. If you want to archive your staged data, in the staging job add another transformation that copies the staged table to a secure location. To do so, you can use the Extract or User-Written Code transformations that are provided by SAS Data Integration Studio in the **Data** section of the Transformations tree. (For information about these transformations, see “Working with Extract Transformations” and “Working with User-Written Code” in the Help for SAS Data Integration Studio.)

Storing Data Tables on z/OS Systems

For best results, store data tables in the UFS file system instead of the traditional z/OS file system.

Note: The documentation for SAS on z/OS has traditionally used the terms UNIX System Services (USS) and hierarchical file system (HFS) to refer to the UNIX file system on z/OS. The SAS 9.3 documentation uses the terms UNIX file system and UFS to refer to this file system. UFS provides a hierarchical storage system (like UNIX or Windows) with outstanding performance. Using UFS for storing data tables can simplify the management of these tables and offer performance benefits over the traditional z/OS file system.

Deleting IT Data Marts

Here are some tips that can help you delete IT data marts effectively:

- Deleting an IT data mart might take a long time, depending on the contents of that IT data mart.
- When the IT data mart was created, the physical location for the Admin library was created. As certain wizards are executed, such as the Adapter Setup wizard, other physical locations are created. These physical locations each contain a `readme.txt` file that carries information about the IT data mart to which the library belongs. When that IT data mart is deleted, the physical tables that belong to that IT data mart are deleted. However, the directory structure is not deleted, and the `readme.txt` files remain for identification purposes. The physical directories that belonged to a deleted IT data mart are empty and can be removed from the location where they were specified.

Tips for Working with Aggregations

Here are some tips that can help you specify ID variables effectively in aggregations:

- ID column suggestions

A useful technique for specifying ID variables is to choose columns that are based on or determined by the value of a column that is already in the class list. An example of this is to choose DATE or DATETIME-based columns that are less granular than the aging column.

For example, consider the case where you have an aggregation that is aged according to the DAYDATE column, which records daily date values. The columns WEEKDATE and MONTHDATE are excellent candidates for use as ID columns because they are (essentially) determined by the value of DAYDATE. Thus, when creating reports, you can easily determine the days that belong with a particular week.

The same is true for other types of date or datetime-based columns that the user might create, such as a column that represents the calendar quarter or accounting period.

- ranking of DATE/DATETIME-based class or ID columns

This tip is a corollary to choosing ID columns based on date, or datetime values. If you have such columns in the class or ID list, then requesting a descending rank on such a column simplifies the report task of reporting on the "last N time periods" (such as the last three weeks). For example, if WEEKDATE is an ID column in an aggregation that has DAYDATE in the class list, then a descending rank on WEEKDATE facilitates reporting on the last three weeks of daily values. The report simply needs to filter where WeekDateRankDescending \geq 3. (All the supplied aggregations and information maps make use of this feature.)

Tips for Working with Information Maps

Using SAS Information Map Studio

You can use SAS Information Map Studio to test existing information maps and to create more complex information maps that join tables, establish pre-filters, and change aggregation statistics. If you do modify an existing information map with SAS Information Map Studio, save the new map under a different name.

Troubleshooting Problems with Information Maps

You receive the following error message "ERROR: Failed to insert filter *filter name*."

This error message is displayed in the SAS log if the user does not have the following privileges: Read Metadata, Write Metadata, Read, and Delete. This error message might also be displayed if the user's password is not registered in SAS Management Console. To register this password, perform the following steps:

1. Start SAS Management Console.
2. Select the **Foundation** repository.
3. Select **User Manager** to display the list of user names that can be accessed on this repository.
4. From the list of names, select the user whose password you want to register.
5. Double-click to open the Properties dialog box for that user.
6. Click **Modify** to open the Edit Login Properties dialog box.
7. Select the **Logins** tab.
8. Enter the password of the user in the **Password** and the **Confirm Password** fields.
9. Click **OK**.
10. Click **OK** to return to the list of user names displayed by the **User Manager**.
11. Close SAS Management Console.

Tips for Installing and Configuring SAS IT Resource Management

Installing and Configuring SAS IT Resource Management

Here are some tips that can help you install and configure SAS IT Resource Management effectively:

- Customizations have been made to the Application Server configuration (sasv9.cfg) and autoexec (appserver_autoexec.sas) files for SAS IT Resource Management. If you need to make additional customizations for site-specific reasons, you should instead modify sasv9_usermods.cfg or appserver_autoexec_usermods.sas.
- The metadata repository for SAS IT Resource Management can become very large, and the SAS Metadata Server also requires substantial memory. Install the physical repository on the fastest media available. Allocate the SAS Metadata Server as much memory as possible.
- Starting and stopping the servers can require several minutes.
- Occasionally, when starting the Object Spawner immediately after starting the SAS Metadata Server, the Object Spawner will fail because it is attempting to contact the SAS Metadata Server. To avoid this situation when automating start-up and shutdown procedures, put sleep or delaying steps into the scripts to allow the SAS Metadata Server to fully initialize before the object spawner attempts to contact it.

Handling Holidays

Overview of Handling Holidays

SAS IT Resource Management 3.2 does not implement any specific holiday support. However, you can use the features of SAS IT Resource Management 3.2 to prevent data that was generated on a holiday from being processed into an aggregation table. To accomplish this goal, perform the following steps.

1. (Optional) If your site's holidays can be specified as individual dates, define a SAS format to identify them. To do this, create a format that defines whether a given date is a holiday. For best results, create this format in the ADMIN library of the IT data mart. (This library is included in the SAS format search concatenation.)

You might want to specify holidays by using different date or datetime measurements (such as 8 AM on a given date). To do so, keep the entire holiday list in the source code of a formula. (This method is explained in the next step.)

Note: It is preferable to use formats to define holidays because formats can be easily revised to reflect changes to the list of holidays, and you do not need to redeploy any jobs. If you use formulas to define your list of holidays, any changes to that list require that jobs be redeployed.

2. Define a formula in the SAS IT Resource Management client that identifies holidays. This formula will be used to create a computed column in the staged table.

To create a formula, select **File** ⇒ **IT Resource Management** ⇒ **Formula**. The New Formula wizard opens. On the General Information page of the wizard, specify a name for the formula. On the Expression page, enter the code to assign the **Y** or **N** value to the computed column. If you created a format in the previous step, use it in this formula.

In some cases, the holiday rule is not a simple assignment statement, based on a format. The rule that identifies whether data is generated on a holiday might require a more complex expression.

Note: If an expression is revised in order to accommodate a change in the list of holidays for a site, you will be required to redeploy the jobs that use that formula.

For information, see [“About Formulas” on page 97](#).


3. (Optional) Specify a distinct value for SHIFT if the data is for a holiday. To do this, you must modify the SHIFT formula to include the code to handle the holiday.

Note: Formulas cannot use other formulas because the order of computation is not guaranteed.

For information about the effects of changing a formula expression, see [“Consequences of Changing a Formula” on page 103](#).

Note: To propagate any change to formulas used in existing jobs, you must redeploy those jobs.

4. Use the holiday formula (and, optionally, the SHIFT formula) to define a computed column in your staged table. This computed column specifies whether the data in this row of the table was generated on a holiday. For information about adding a computed column to a staged table, see [“Modify a Staged Table” on page 122](#).
5. Define an input filter that prevents holiday data from being processed into the aggregation tables. You can define filters for the aggregation table in the Aggregation transformation that generates the table. To do so, perform the following steps:
 - a. Double-click the job that contains the Aggregation transformation. The job opens in the **Diagram** tab of the Job Editor window.
 - b. Right-click the Aggregation transformation. From the Properties dialog box, select the **ITRM Options** tab. On this tab, you can define a filter by clicking **New**. In the List of Filters box, you can add a filter that can be used for this Aggregation transformation.

Display A10.1 ITRM Options Tab

Note: Set the filter so that data that satisfies its conditions is the only data that is processed.

After you specify the holiday filter, click **OK** and save the job.

6. After the filter is added to the Aggregation transformation, you can use it with one or more aggregation tables that are generated by this transformation. To apply a filter to an aggregation table, do one of the following two steps.

- For a new aggregation table that you want to create, right-click the Aggregation transformation and select **Add Aggregation Table** from the Properties dialog box. Select the type of table that you want to create (summarized or simple.)

Respond to the prompts of the wizard that opens. On the second page of the wizard, you can select the holiday filter that you want to use. When you have completed entering the information for the table, click **Finish** and save the job.

- For an existing aggregation table, right-click the table and select **Edit Summarized Aggregation Table** or **Edit Simple Aggregation Table** from the Properties dialog box. Then navigate to the Apply Filter page of the wizard and select the holiday filter that you want to use. When you have completed editing the table, click **Finish** and save the job.

Note: Implementing a filter does not change existing data in a table.

Note: Newly created jobs or jobs that have been changed must be saved and redeployed. For information about redeploying jobs, see [“Redeploy All Jobs on the Server” on page 281](#).

Example of Handling Holidays

Overview of the Example

In this example, a format is specified that defines Christmas Day as a specific holiday. A formula is created that uses the holiday format to specify a Holiday column. The Holiday column is added to a staged table as a computed column. This computed column will contain a **Y** for every row of data whose date column corresponds to the definition of the holiday.

After the staged table is populated with the computed Holiday column, a filter is created in the Aggregation transformation. The filter prevents holiday data from being processed into the aggregation table. This filter can be implemented for any aggregation table that is created or updated by that Aggregation transformation.

Creating the Format

The following code creates a SAS format that specifies the Christmas holiday for 2010 and 2011. The resulting value of **holiday** is a single-character of **'Y'** or **'N'** that indicates whether the date of the data is a holiday.

```
libname ADMIN BASE "c:\datamarts\smfasw\admin";
proc format library=admin;
  value holiday (default=1 min=1 max=1)
    '25Dec2010'd = 'Y'
    '25Dec2011'd = 'Y'
    other = 'N';
run;
```

Creating the Formula

The following formula uses the value of the **holiday** format to define the Holiday column.

```
Holiday = PUT(DAYDATE,holiday.);
```

Note: Add the Holiday column to the staged table as a computed column. This column can be used as a filter to exclude from the Aggregation transformation all data that was generated on a holiday. The filter ensures that only data that is not a holiday will be processed.

(Optional) Including the Holiday Specification in the SHIFT Formula

Suppose your SHIFT formula currently has the following expression:

```
if weekday(datepart(datetime)) in (1,7) then shift = '3';
else if timepart(datetime) < '08:00:00't or
    timepart(datetime) >= '17:00:00't then shift = '2';
else shift = '1';
```

In order to accommodate the holiday specification, change this expression as follows:

```
if put(DAYDATE,holiday.) eq 'Y' then shift = '0';
else if weekday(datepart(datetime)) in (1,7) then shift = '3' ;
```

```

else if timepart(datetime) lt '08:00:00't or
      timepart(datetime) ge '17:00:00't then shift = '2';    else shift = '1';


```

Define an Input Filter for Holiday Data in an Aggregation Transformation

Open the job that contains the appropriate Aggregation transformation. Right-click the transformation. On the Properties dialog box of the Aggregation transformation, select the **ITRM Options** tab. Click **New** to define the following filter. In the following display, the Holiday filter has this expression:

Holiday eq 'N'

Display A10.2 ITRM Options Tab

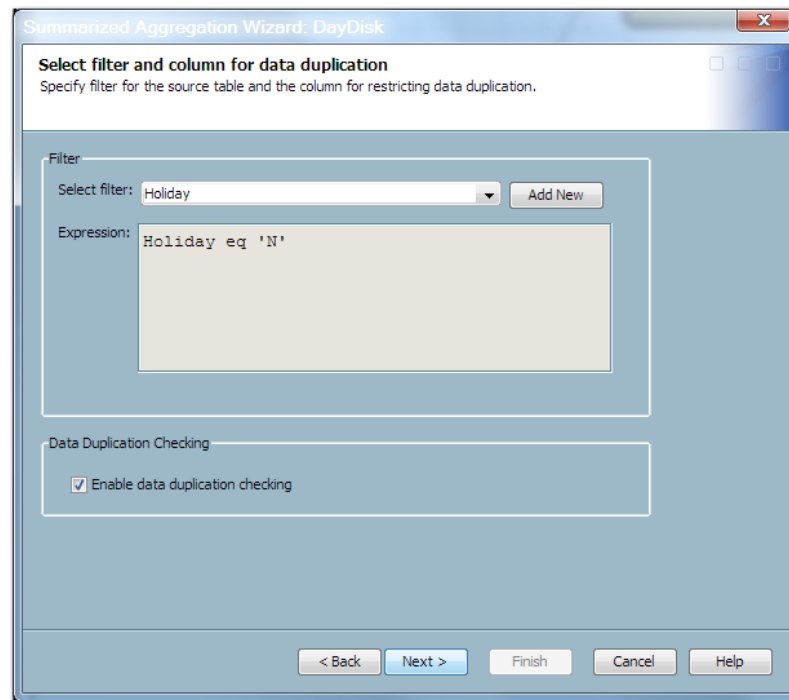


Click **OK**.

Use the Holiday Filter with One or More Aggregation Tables

The Aggregation wizards enable you to specify the Holiday filter for your aggregation table. The following display shows the page of the Aggregation wizards where you can select the filter.

Note: The filter is an inclusive filter. Data that satisfies the conditions of the filter is processed into the aggregation table.

Display A10.3 *Selecting a Holiday Filter*

After you select the input filter, save the job and redeploy it for execution.

Note: The input filter does not affect data that is already in the aggregation table. However, it will filter any new incoming data to the aggregation table.

Performance Issues

Improving Memory Performance

Out-of-memory issues can cause ETL jobs to fail.

To resolve out-of-memory issues on UNIX platforms, set the following SAS options as shown:

- **MEMSIZE=512M**
- **SORTSIZE=256M**
- **BUFSIZE=48K**
- **BUFNO=1500**

To resolve out-of-memory issues on z/OS platforms, set the following SAS options as shown:

- **SORTSIZE=256M**
- **BUFSIZE=48K**
- **BUFNO=1500**

Note: In addition, using SAS IT Resource Management software on z/OS requires that batch jobs have a REGION of at least 1024M. A parallel ASSIZEMAX of 1024M is also required.

To resolve out-of-memory issues on Windows (64-bit) platforms, set the following SAS options as shown:

- **MEMSIZE=2G**
- **SORTSIZE=256M**

TIP SAS options can be set at configuration time, invocation time, or run time.

Note: If you continue to experience out-of-memory issues, you might need to further increase the value of the MEMSIZE option. For information about metadata memory issues, see [“Metadata Server Memory Issues” on page 31](#).

Limiting the Processing of SMF Records by System

An aggregation filter can be used to limit the SMF records that are processed into an aggregation table to specific systems. However, the staging code must process the data that will not be aggregated. In addition, the filter must be set on each SMF aggregation. It is more efficient to filter data during the staging job processing.

To limit the SMF records to process only specific systems during a staging job process, add an IF statement in the IMACFILE MXG SOURCLIB member. For example, if you want to process only those records from the 'PROD' and 'TEST' systems, place a copy of the IMACFILE MXG SOURCLIB member into your MXGUSER source library.

Modify the user IMACFILE MXGUSER member with the following SAS statement:

```
IF SYSTEM IN ('PROD','TEST'); /*Only process records for these systems */
```

where the systems for which you want to process SMF records are listed in the IN clause. For additional examples or information about coding the IF statement, please refer to the documentation found in the IMACFILE MXG SOURCLIB member.

Note: Make sure you have a working backup copy of your SAS IT Resource Management installation files and IT data marts before you make any changes.

Staging Jobs Fail Due to Lack of Space

Staging job fails due to lack of space on the WORK library on Windows.

The location of the Work library defaults to the user's C: drive. The space on this drive might not be adequate for running a staging job. To resolve this issue, edit the **sas.cfg** file to point WORK to a location that has more space.

Aggregation Jobs Require Lengthy Processing Times

Some aggregation jobs that contain rank variables or join columns can require a long time to finish.

For best results, minimize the number of rank variables or join columns that you specify in a summarized aggregation table.

Information Map Jobs Require Lengthy Execution Times

Executing a job that contains multiple information maps can take a long time. Information maps facilitate the process of generating reports on your IT data and are the required input to SAS Web Report Studio. Jobs that generate many information maps can take a long time to execute. Although SAS Enterprise Guide

can generate reports from tables or information maps, using information maps is preferred.

The Adapter Setup wizard creates separate jobs for Information Map transformations. You do not need to schedule these Information Map jobs to run nightly. These Information Map jobs need to be run initially only once—after the underlying tables are written, and later if something changes about the structure of the tables for which information maps are to be created. For example, if you change the class columns or statistics in an Aggregation transformation, that will change the resulting aggregation table. In that situation, the corresponding Information Map must be recreated as well.

If you are sure that you will never use specific information maps, remove the unneeded Information Map transformation from the Process Flow Diagram before you deploy the job.

Appendix 11

Statistics

SAS IT Resource Management can calculate the following statistics:

- count
- geometric mean
- harmonic mean
- maximum
- mean
- minimum
- number missing
- population coefficient of variance
- population standard deviation
- population variance
- range
- sample coefficient of variance
- sample standard deviation
- sample variance
- sum
- uncorrected sum of squares
- weighted geometric mean
- weighted harmonic mean
- weighted mean
- weighted population coefficient of variance
- weighted population standard deviation
- weighted population variance
- weighted sample coefficient of variance
- weighted sample standard deviation
- weighted sample variance
- weighted sum
- weighted uncorrected sum of squares

SAS IT Resource Management can also calculate the period-to-period percent change of a statistic within class combination.

Index

Special Characters

%RMDELETE macro
 options [512](#)
 overview [511](#)
 required arguments [511](#)
 syntax [511](#)
 %RMDELETE macro [511](#)
 %RMDUPCHK macro [513](#)
 example [516](#)
 examples [512](#)
 notes [512](#), [516](#)
 options [515](#)
 required arguments [514](#)
 syntax [513](#)
 %RMDUPDSN macro [516](#)
 example [517](#)
 notes [517](#)
 required arguments [517](#)
 syntax [517](#)
 %RMDUPINT macro [517](#)
 example [518](#)
 notes [517](#)
 syntax [517](#)
 %RMDUPUPD macro [518](#)
 example [518](#)
 notes [518](#)
 syntax [518](#)
 %RMMSBMCP macro [519](#)
 options [519](#)
 required arguments [519](#)
 syntax [519](#)
 %RMMSDTPS macro [521](#)
 options [521](#)
 required arguments [521](#)
 syntax [521](#)
 %RMSSSNMP macro [523](#)
 options [523](#)
 required arguments [523](#)
 syntax [523](#)
 %RMPOINT macro [524](#)
 example [529](#)

 notes [525](#)
 options [525](#)
 required arguments [525](#)
 syntax [525](#)
 %RMRUNETL macro [529](#)
 options [530](#)
 required arguments [530](#)
 running an overall job [282](#)
 syntax [530](#)

A

access command [356](#)
 accessibility [14](#)
 adapter and domain category
 specifications [470](#)
 Adapter Setup wizard [43](#)
 about [177](#)
 accessing and navigating [179](#)
 advantages of [178](#)
 aggregation jobs and objects created by
 [196](#)
 creating jobs [180](#)
 creating metadata for aggregation tables
 [133](#)
 deploying and running jobs created by
 [205](#)
 determining if reporting jobs will be
 created by [200](#)
 domain categories [204](#)
 generating jobs [264](#)
 information map filters created by [220](#)
 information map jobs and objects
 created by [199](#)
 information map objects created by [219](#)
 IT data processing [264](#)
 jobs created for domain categories [381](#)
 jobs types created by [191](#)
 libraries created by [202](#)
 log file created by [203](#)
 naming conventions [178](#)

- naming information map filters 221
 - naming information map objects 219
 - purpose of 4
 - renamed Quick Start wizard 13
 - reporting jobs and objects created by 199
 - staging jobs and objects created by 195
 - steps to complete 179
 - storing generated objects 193, 265
 - storing information map objects 220
 - adapter-based subfolders 54
 - adapters 13, 67
 - external names 127
 - MXG based, filerefs for 283
 - MXG based, for duplicate-data checking 396
 - MXG views 330
 - non MXG, for duplicate-data checking 399
 - staging data 75
 - staging parameters for 362
 - summary of 467
 - supported 290
 - template tables 75
 - types of 68, 291
 - Admin library 53
 - age limit, for data acquisition table 356
 - aggregating staged tables 9
 - aggregation
 - concepts 461
 - jobs created by Adapter Setup wizard 196
 - strategies 463
 - table groups and time periods 462
 - wizards 44, 133
 - aggregation table wizards 44
 - aggregation tables 132
 - computed columns in 126
 - creating metadata with Adapter Setup wizard 133
 - creating metadata with aggregation wizard 133
 - creating with New Aggregation Table from Templates wizard 159
 - deleting filters from 167
 - deleting from Aggregation transformation 160
 - deleting from IT data marts 161
 - erasing 161
 - generating physical aggregation table 134
 - indexing 173
 - libraries for 203
 - modifying 160, 162
 - properties 134
 - purging contents 164
 - purging data from process flow diagram 164
 - purging data on IT Data Marts tab 164
 - renaming 164
 - setting up Aggregation transformation in process flow diagram 140
 - simple 132
 - summarized 132
 - Aggregation transformation 9, 13
 - Aggregation transformation filters 165
 - accessing ITRM Options tab 165
 - defining 166
 - deleting 166
 - Expression Builder and 168
 - modifying 167
 - Aggregation transformations 9, 170
 - creating metadata with aggregation wizard 133
 - deleting aggregation tables from 160
 - deleting from jobs 173
 - properties 170
 - renaming 173
 - setting up in process flow diagram 140
 - aggregations
 - key metrics 417
 - naming conventions 503
 - tips for working with 542
 - aggregations by time period 196
 - Aggregations through Template 44
 - aging
 - options for summarized aggregation tables 142
 - specifying criteria 143
 - analysis variables
 - creating statistics for 146
 - specifying for summarized aggregation tables 145
 - applications data sources 70, 292
 - architecture 6
 - ASG TMON2CIC adapter, customizations 333
 - ASG TMONC2CIC staging parameters 362
 - ASG TMONDB2 staging parameters 363
 - ASUMCICX
 - customizations for the CICS domain category 332
 - customizations for the DB2 domain category 332
- B**
- backing up SAS Metadata Repository 31
 - backup process 536
 - backups
 - data 537

- metadata 537
- process 536
- reports 538
- SAS Metadata Repository 29
- Base SAS sessions
 - browsing metadata repository from 30
- Basic Properties pane 40
- batch jobs 276
- batch maintain table command
 - examples 93
 - for UNIX 94
 - for Windows 94
 - for z/OS 95
 - options 91
 - syntax 91
- batch mode
 - deploying jobs for 18
- BMC Mainview IMS staging parameters 364
- BMC Perf Mgr
 - extracting raw data 304
 - preparing raw data 303
 - raw data 303
 - staging parameters 365
 - Windows domain category 472

C

- CA TMS
 - staging parameters 365
- CA TMS adapter
 - customizations for the Jobs domain category 333
- character data, gathering with snmpwalk 361
- Checkouts tree 40
- class variables, for summarized aggregation tables 144
- client
 - See SAS IT Resource Management client
- client tier 6
- client tier components 288
- collecting raw data 8
- collectors 13
- columns
 - See also computed columns
 - computed 153
 - deleting from indexes 175
 - deleting from staged tables 124
 - external names of columns in staged tables 127
 - importing into staged tables 123
 - in staged tables 124
 - in template tables 79
 - labeling 460

- modifying properties in staged tables 123
- removing join columns 152
- specifying column detail for summarized aggregation tables 145
- specifying join columns 151
- computed columns 97, 153
 - adding 153
 - adding to a table 99
 - adding to staged tables 122
 - deleting 156
 - in aggregation tables 126
 - in data model 460
 - in staged tables 126
 - in template tables 82
 - modifying 99, 154
- configuration 544
- Connection Profile dialog box 35
- Connection Profile wizard 35
- connection profiles 35
- control data sets 392
 - intermediate 392
 - permanent 393
 - temporary 392
- counters, fast-moving 527
- custom tables 328

D

- data, backing up 537
- data acquisition table
 - age limit for 356
- data administrators
 - architecture and basic principles for 6
- data base adapters
 - accessing raw data 315
- data collection sets (DCS)
 - for ITRM exchange 303
 - for ITRM Windows server 302
- data columns
 - adding to staged tables 122
- data marts
 - See also IT data marts
 - archiving 541
 - deleting 58, 542
 - managing 541
 - storing 542
 - tips for working with 541
- data model 459
 - adapter and domain category specifications 470
 - adapter and domain category summary 467
 - aggregation key and ranked metrics 462
 - aggregation strategies 463

- aggregation table groups and time periods 462
- features 460
- staging and aggregation concepts 461
- data sets, control 392
- data sources 68, 291
 - supported adapters 290
 - user-written 74, 297
- database adapters
 - duplicate-data checking 316
 - preparing for staging 315
- date values
 - National Language Support (NLS) 12
- datetime values
 - National Language Support (NLS) 12
- debugging 539
 - Java objects for 540
- default duration 357
- delimiter in raw data 357
- deploying jobs 274
 - adding locations for 18
 - allocating raw data file externally 282
 - created by Adapter Setup wizard 205
 - redeploying 280
 - redeploying all jobs on server 281
- desktop 36
- detail aggregated tables 198
- detail aggregation tables 196
- detail tables 9
- detail-level data 13
- documentation 10
- domain categories 204
 - jobs created by Adapter Setup wizard 381
 - specifications 470
 - summary 467
- DT Perf Sentry
 - accessing ITRM exchange DCS 303
 - accessing ITRM Windows server DCS 302
 - accessing raw data 300
 - Exchange domain category 471
 - extracting raw data 301
 - preparing raw data for staging 300
 - Server domain category 470
 - staging parameters 366
- DT Perf Sentry with MXG
 - staging parameters 367
- duplicate-data checking 357, 392
 - control data sets for 392
 - examples 414
 - for non MXG user-written staging code 410
 - for SAP ERP adapter 414
 - implementing 396
 - implementing macros for 394

- macros for 394
- overview 391
- supplied MXG based adapters 396
- supplied non MXG adapters 399
- user-written staging code based on MXG code 403
- duplicate-data report 392
- duration
 - default 357
- duration, presummarization 360

E

- existing users, adding 17
- Expression Builder 168
 - accessing 99
 - IT Formulas tab 98
- external names
 - changing for table or variable 128
 - for MXG adapters 129
 - of staged tables and columns 127
 - per adapter 127
 - viewing for columns 129
 - viewing for tables 128

F

- filerefs for MXG based adapters 283
- filters
 - Aggregation transformation filters 165
 - deleting from aggregation tables 167
 - information map filters 220, 221
 - selecting for summarized aggregation tables 142
- Folders tree 38
- formats
 - National Language Support (NLS) 12
- formulas 97
 - accessing 98
 - consequences of changing 103
 - copying 101
 - creating 98, 100
 - creating with New Formula wizard 100
 - deleting 102
 - in data model 461
 - modifying 102
 - renaming 104
 - usage 97
- FTP 358
- future data 359

G

- galleries 4
- Gallery Manager 4
- groups of users, adding 16

H

- Help 12
- HP Perf Agent
 - domain categories 474
 - raw data 298
 - staging parameters 368
- HP Reporter
 - domain categories 474
 - staging parameters 368
- HP Reporter adapter
 - duplicate-data checking for 414

I

- IBM DCOLLECT
 - staging parameters 369
- IBM EREP
 - staging parameters 370
- IBM SMF
 - staging parameters 371
- IBM SMF custom tables 328
- IBM TPF
 - staging parameters 372
- IBM VMMON
 - staging parameters 373
- icons, for objects 46
- ID variables
 - allowing duplicates 356
 - for summarized aggregation tables 144
- IMACEXCL
 - customizations for the CICS domain category 332
 - customizations for the DB2 domain category 332
- IMACINTV
 - customizations for Jobs domain category 332
- IMACOWT
 - customizations for the CICS domain category 333
 - customizations for the DB2 domain category 333
- IMACSHFT
 - customizations for RMF 331
- IMACSPIN
 - customizations for Jobs domain category 331
- IMACUOW
 - customizations for the CICS domain category 332
 - customizations for the DB2 domain category 332
- IMACWORK
 - customizations for RMF 331
- importing columns into a staged table 123
- indexes 137, 173
 - accessing 174
 - creating 175
 - deleting 176
 - deleting columns from 175
 - indexing aggregation tables 173
 - renaming 176
- information map filters 221
 - accessing 222
 - adding 224
 - benefits of 221
 - created by Adapter Setup wizard 220
 - deleting 225
 - modifying 225
 - naming 221
- information map jobs, created by Adapter Setup wizard 199
- information map objects
 - generated by Adapter Setup wizard 219
 - naming 219
 - storage by Adapter Setup wizard 220
- Information Map Parameters tab 222
- Information Map transformations 210
 - accessing properties of 211
 - creating 212
 - deleting from jobs 217
 - errors running 230
 - modifying 215
 - properties 210
 - specifying 212
 - user-specified 212
- information maps 208
 - accessing in SAS Information Map Studio 227
 - benefits of using 208
 - creating 208
 - creating with SAS Information Map Studio 228
 - deleting 218
 - deleting items in 228
 - error opening or accessing data from 230
 - generating 9
 - lengthy processing time 230
 - modifying items in 228
 - naming conventions 505
 - not appearing in IT Data Mart folder 229
 - permissions for 16
 - renaming items in 228
 - SAS Information Map Studio 226, 543
 - storage of 209
 - testing 229
 - tips for working with 543
 - troubleshooting 229
 - updating from a template 218
- installation and configuration 544

- insufficient disk space 129
 - intermediate control data sets 392
 - intermediate staging view 361
 - Inventory tree 39
 - IT data
 - jobs for processing 263
 - processing with Adapter Setup wizard 264
 - processing without Adapter Setup wizard 267
 - IT Data Mart folder
 - information maps not appearing in 229
 - IT data mart objects 53
 - IT data marts 13, 51
 - accessing 52
 - contents of 53
 - creating 52, 57
 - deleting 58
 - deleting aggregation tables from 161
 - erasing 58
 - exporting 60, 61
 - importing 60, 62
 - modifying 59
 - New IT Data Mart wizard 52
 - populating 52
 - promoting 60
 - purging tables of 59
 - renaming 60
 - storage location 52
 - IT Data Marts tab
 - purging data from aggregation tables 164
 - IT Data Marts tree 38
 - deleting tables from 271
 - deleting tables with contents 271
 - purging tables from 272
 - IT Formulas tab 98
 - IT Resource Management (ITRM)
 - using 8
 - ITRM
 - See also* SAS IT Resource Management (ITRM)
 - permissions for 18
 - ITRM Options tab
 - Properties dialog box 165
 - itrmusers group 16
- J**
- JES 360
 - job deployment, adding locations for 18
 - Job Editor 40
 - job failure, recovering 536
 - job properties 41
 - jobs
 - See also* Adapter Setup wizard
 - adding an existed staged table to 120
 - adding transformations to 269
 - created by Adapter Setup wizard 191
 - created by Adapter Setup wizard, for domain categories 381
 - creating 268
 - creating with Adapter Setup wizard 180
 - deleting 269
 - deleting Aggregation transformations from 173
 - deleting Information Map transformations from 217
 - deleting staged tables from 121
 - deleting transformations from 270
 - deleting User-Written Staging transformation from 258
 - deploying and running jobs created by Adapter Setup wizard 205
 - deploying for scheduling 274
 - generating with Adapter Setup wizard 264
 - modifying 269
 - processing IT data 263
 - redeploying 280
 - redeploying all jobs on server 281
 - renaming 273
 - running 273
 - running an overall job with %MRUMETL macro 282
 - running immediately 273
 - scheduling to run in batch 276
 - scheduling with built-in or third-party scheduler 276
 - scheduling with Schedule Manager 277
- join columns
- removing 152
 - specifying for summarized aggregation tables 151
- K**
- key metrics 417
 - aggregations 196, 197, 417
 - in data model 462
 - network performance 419, 429
 - reporting areas 418
 - resource utilization 419
 - server performance 419
 - SMF adapter 432
 - storage usage 426
 - Web Log performance 426
 - workload profiling and characterization 419, 427

L

libraries created by Adapter Setup wizard 202

local time, converting VMware data 126

locale 13

locations, adding for job deployment 18

log files

- created by Adapter Setup wizard 203

logical record length 360

logon

- to metadata server 34

M

machine name 360

macros 510

- %RMDELETE 511
- %RMDUPCHK 513
- %RMDUPDSN 516
- %RMDUPINT 517
- %RMDUPUPD 518
- %RMMSBMCP 519
- %RMMSDTPS 521
- %RMMSSNMP 523
- %RMPOINT 524
- %RMRUNETL 529

for duplicate-data checking 394

process macros 13

reduction macros 13

Maintain Staged Tables wizard

- about 84
- using across all IT data marts 90
- using to update tables within an IT data mart 85
- wizards 44
- working with 84

maintaining staged tables

- about 90
- in batch mode 90

management information base (MIB)

- definitions 314

memory

- client and server issues 31
- SAS IT Resource Management metadata server 31

memory issues

- client 31
- server 31

metabrowse command 31

metadata 27

- accessing ITRM metadata 30
- backing up 537
- backing up and restoring SAS Metadata Repository 29
- browsing repository from Base SAS sessions 30

- creating for aggregation tables 132
- creating with Adapter Setup wizard 133
- creating with aggregation wizard 133
- exploring SAS Data Integration Studio metadata repository 30
- generating for transformations 267
- supplied 28
- system authentication when logging on to SAS Metadata Repository 28

metadata repository

- See also* SAS Metadata Repository
- from Base SAS sessions 30
- SAS Data Integration Studio 30

metadata server

- See also* SAS Metadata Server
- logging on to 28, 34
- memory issues 31

Microsoft Office products 5, 10

middle tier 6

middle tier components 287

migration wizard 22

- prerequisites 22
- running 22

MIPS and MSU columns

- calculation dependencies 336
- calculations 343
- including in staged tables 336
- initializing 337
- preparing 337
- staging methodology 339

MS SCOM

- Exchange domain category 499
- Server-Based domain categories 498
- staging parameters 374

MS SCOM adapter

- duplicate-data checking for 414

multi-tier processing 6

MXG

- accessing raw data 298
- code 330
- extracting raw data from 298
- views 330

MXG adapters

- external names for 129
- filerefs for 283

MXG based adapters

- for duplicate-data checking 396

N

names

- external names for MXG adapters 129
- external names of staged tables and columns 127

naming conventions

- Adapter Setup wizard 178

- aggregations 503
- information map filters 221
- information map objects 219
- information maps 505
- ITRM objects 503
- object locations and 503
- reporting 507
- staging 508
- National Language Support (NLS) 12
- network performance metrics 429
- New Aggregation Table from Templates wizard
 - creating aggregation tables 159
- New Formula wizard 43
 - creating formulas 100
- New IT Data Mart wizard 44, 52
 - creating IT data marts 52
 - invoking 52
- new users, adding 17

O

- objects
 - created by Adapter Setup wizard 195, 196
 - icons for 46
 - naming conventions and location of 503
 - naming conventions for 503
- online Help 12

P

- PDB 13
- percent change
 - for a statistic 149
 - specifying for summarized aggregation tables 149
- permanent control data sets 393
- permissions
 - for information maps 16
 - setting 18
- preliminary tasks 7
- presummarization duration 360
- process flow diagrams
 - deleting tables from transformations 270
 - deleting tables with contents 271
 - purging data from aggregation tables 164
 - purging tables from transformations 272
 - setting up Aggregation transformation in 140
- process macros 13
- processing raw data 8

properties

- aggregation tables 134
- Aggregation transformations 170
- columns in staged tables 123
- Information Map transformations 210, 211
- job properties 41
- staged tables 113
- staging transformations 108
- table properties 42
- User-Written Staging transformation 234
- Properties dialog box 39
 - ITRM Options tab 165
- purging
 - contents of aggregation tables 164
 - data from aggregation tables, on IT Data Marts tab 164
 - data from aggregation tables from process flow diagram 164
 - options for summarized aggregation tables 142
 - specifying criteria 143
 - tables from IT Data Marts tree 272
 - tables from transformations in process flow diagrams 272

Q

- QuickStart wizard 13

R

- ranked metrics 462
 - SMF adapter 432
- ranking
 - deleting rank variables 151
 - for summarized aggregation tables 150
 - specifying for a variable 150
- raw data
 - adapters and 67
 - allocating files externally 282
 - collecting 8
 - delimiter in 357
 - input directory 360
 - input file 361
 - input library 361
 - input type 357
 - processing (staging) 8
 - staging 106
 - type 356
- redeploying jobs 280
- reduction macros 13
- renaming
 - aggregation tables 164
 - Aggregation transformations 173

- formulas 104
- indexes 176
- IT data marts 60
- items in information maps 228
- jobs 273
- User-Written Staging transformation 259
- reporting
 - key metrics and 418
 - naming conventions 507
 - Web-based 10
- reporting jobs
 - created by Adapter Setup wizard 199
- reports
 - backing up 538
 - defining and generating for summarized data 10
 - viewing 10
- resource utilization metrics 419
- resources
 - collecting raw data about 8
- restoring SAS Metadata Repository 31
- RMFINTRV
 - customizations for RMF 331
- rsh host command 361
- running jobs 273
 - immediately 273
 - in batch 276
 - running an overall job with %RMRUNETL macro 282
- Runtime Manager pane 41

S

- SAP ERP
 - staging parameters 374
- SAP ERP adapter
 - accessing raw data 323
 - duplicate-data checking for 414
 - preparing raw data 323
- SAPR3
 - Batch Job and Workload domain categories 477
 - Systems, Machines, Tasks, and Transactions domain categories 475
- SAR
 - accessing raw data 306
 - extracting raw data 306
 - preparing raw data 306
 - SAR domain category 478
 - staging parameters 374
- sar command
 - using 307
- SAS Add-In for Microsoft Office 5, 10
- SAS Business Intelligence Platform 5

- SAS Data Integration Studio 4, 33, 74, 297
 - configuring data extraction jobs 326
 - exploring metadata repository 30
 - metadata 27
 - tips for working with 541
 - with plug-ins for ITRM 27
 - wizards supplied by 45
- SAS Enterprise Guide 5
 - accessing SAS Metadata Repository 20
 - metadata 27
 - migrating projects 22
 - setting up for ITRM data 20
 - working with 20
- SAS Information Delivery Portal 5
- SAS Information Map Studio 5, 226, 543
 - accessing 227
 - accessing information maps in 227
 - benefits of 227
 - creating information maps 228
 - modifying, renaming, or deleting items in information maps 228
 - overview 226
 - testing information maps 229
- SAS IT Resource Management
 - data model 459
 - documentation 10
 - logging on to SAS Metadata Server 28
 - metadata server memory issues 31
 - tips for installing and configuring 544
- SAS IT Resource Management (ITRM) 3
 - accessibility 14
 - architecture 6
 - basic principles and components 7
 - components 287
 - features 3
 - functionality 4
 - preliminary tasks 7
 - starting 8, 34
 - terminology changes 13
 - wizards supplied by 43
- SAS IT Resource Management client 33
 - features 36
 - logging on to metadata server 34
- SAS Management Console 15
 - accessing 16, 317
 - adding locations for job deployment 18
 - adding users and groups of users 16
 - creating SAS metadata objects 325
 - introduction to 15
 - metadata 27
 - Schedule Manager 277
 - working with 15
- SAS Metadata Repository 5
 - accessing for SAS Enterprise Guide 20
 - backing up and restoring 29, 31

- from Base SAS sessions 30
 - SAS Data Integration Studio 30
 - system authentication when logging on to 28
- SAS Metadata Server
 - logging on to 28
 - logging on to from SAS IT Resource Management 28
 - memory issues 31
- SAS Statistics and Econometric Time Series 5
- SAS Web Report Studio 5, 10
 - metadata 27
- Schedule Manager
 - scheduling jobs 277
- scheduling
 - deploying jobs for 274
 - jobs to run in batch 276
 - with built-in or third-party scheduler 276
- scheduling jobs
 - with Schedule Manager 277
- Section 508 14
- server
 - performance metrics 419
- server performance metrics 419
- server tier 6
- server tier components 287
- servers
 - redeploying all jobs on 281
- simple aggregation tables 44, 132
 - creating with Simple Aggregation Table wizard 157
 - data 13
 - modifying 162
- Simple Aggregation wizard 44
- simple aggregations 9
- site name 361
- SMF
 - adapter metrics 419
 - CICS domain category 487
 - DB2 domain category 490
 - Jobs, TSO and OMVS domain categories 483
 - limiting records 336
 - RMF domain category 479
- SMF adapter
 - batch job, TSO and OMVS session performance domain metrics 449
 - CICS transaction activity domain metrics 452
 - CPU processor domain metrics 448
 - DB2 thread accounting domain metrics 455
 - device domain metrics 446
 - goal mode service domain metrics 438
 - Lpar domain metrics 435
 - paging and swapping domain metrics 447
 - paging file domain metrics 448
 - system domain metrics 432
 - workload domain metrics 441
 - workload period domain metrics 443
- SNMP
 - accessing raw data 308
 - device type lookup table 313
 - extracting raw data 310, 311
 - Network Activity, Device and Interface Data, and System Data domain categories 496
 - sources of raw data 308
 - staging parameters 375
 - using HP NNM tools to extract data 310
 - using round-robin tools to extract data 311
- snmpwalk 361
- spin libraries 203
 - specifying for ASG TMON2CIC 334
 - specifying for IBM SMF 334
- Staged Table wizard 44
- staged tables 112
 - adding a new computed column to 122
 - adding a new data column to 122
 - adding an existing table to a job 120
 - aggregating 9
 - columns in 124
 - computed columns in 126
 - creating 114
 - creating from template tables 83
 - deleting columns in 124
 - deleting from a tree view 121
 - deleting from jobs 121
 - erasing 120
 - external names of 127
 - IBM SMF custom tables and 328
 - importing column from another table into 123
 - library for 203
 - modifying 122
 - modifying column properties in 123
 - properties of 113
 - purging content of 122
 - task list for 114
 - versus template tables 77
- staging
 - naming conventions 508
 - raw data 8
- staging and aggregation concepts 461
- staging code, user-written 70, 293
 - See also *User-Written Staging transformation*

- staging data 106
 - columns in staged tables 124
 - computed columns in staged tables 126
 - external names of staged tables and columns 127
 - insufficient disk space 129
 - methods for 106
 - reasons for 106
 - staged tables 112
 - staging transformations 107
 - VMware 126
 - with MXG code 330
 - staging jobs
 - adding staging transformations to 109
 - created by Adapter Setup wizard 195
 - staging parameters 356
 - ASG TMONC2CIC adapter 362
 - ASG TMONDB2 adapter 363
 - BMC Mainview IMS adapter 364
 - BMC Perf Mgr adapter 365
 - CA TMS adapter 365
 - DT Perf Sentry adapter 366
 - DT Perf Sentry with MXG adapter 367
 - for supported adapters 362
 - HP Perf Agent adapter 368
 - HP Reporter adapter 368
 - IBM DCOLLECT adapter 369
 - IBM EREP adapter 370
 - IBM SMF adapter 371
 - IBM TPF adapter 372
 - IBM VMMON adapter 373
 - MS SCOM adapter 374
 - SAP ERP adapter 374
 - SAR adapter 374
 - SNMP adapter 375
 - user-written 362, 376
 - VMware Data Acquisition adapter 379
 - VMware vCenter adapter 379
 - Web Log adapter 380
 - staging transformations 13, 107
 - adapters and 75
 - adding to a staging job 109
 - deleting 112
 - editing 112
 - properties of 108
 - starting SAS IT Resource Management 8, 34
 - statistics 553
 - creating for analysis variables 146
 - deleting 148
 - percent change calculation for 149
 - removing percent change calculation for 149
 - specifying for summarized aggregation tables 146
 - storage usage metrics 426
 - storing data marts 542
 - summarized aggregation data 13
 - Summarized Aggregation Table wizard
 - creating summarized aggregation tables 141
 - summarized aggregation tables 44, 132
 - aging and purging options 142
 - completing the specification of 157
 - computed columns 153
 - creating with Summarized Aggregation Table wizard 141
 - deleting multiple columns 148
 - entering general information about 141
 - modifying 162
 - ranking 150
 - removing join columns 152
 - selecting filters 142
 - specifying analysis variables 145
 - specifying class and ID variables 144
 - specifying column details 145
 - specifying join columns 151
 - specifying library information 141
 - specifying percent change 149
 - specifying statistics 146
 - summarized aggregations 9
 - summarized data
 - defining and generating reports for 10
 - summary level data 13
 - supplied formulas 97
- ## T
- table groups 462
 - table properties 42
 - tables
 - See also* aggregation tables
 - adding to transformations 270
 - deleting, with contents, from IT Data Marts tree 271
 - deleting from transformation in process flow diagram 271
 - erasing from transformations 272
 - purging from IT Data Marts tree 272
 - purging from transformations in process flow diagrams 272
 - purging IT data mart tables 59
 - staged tables 112
 - template tables 67, 75
 - adapters and 75
 - columns in 79
 - computed columns in 82
 - ITMS properties for 81
 - locating 76
 - properties 77
 - using 83
 - versus staged tables 77

- viewing 77
- temporary control data sets 392
- terminology changes 13
- testing
 - information maps 229
- time periods 462
- timestamps 67
- transformations 8, 45
 - adding tables to 270
 - adding to jobs 269
 - Aggregation transformations 9
 - deleting from jobs 270
 - erasing tables from 272
 - generating metadata for 267
 - purging tables from process flow diagram 272
- Transformations tree 39
- Tree views 38
 - deleting staged tables from 121
- trees 38
- troubleshooting
 - information maps 229

U

- UNIX system data sources 70, 292
- user groups, adding 16
- user-defined formulas 97
 - creating 98
- user-specified Information Map
 - transformations 212
- user-written data sources 74, 297
- user-written staging code 70, 293
 - based on MXG code 403
 - duplicate-data checking for non MXG code 410
- user-written staging parameters 362
- User-Written Staging transformation 233
 - additional generated code 259
 - Code tab 235
 - creating 252
 - deleting from a job 258
 - Extended Attributes tab 235
 - General tab 234
 - modifying 259
 - Notes tab 235
 - Precode and Postcode tab 235
 - properties 234
 - renaming 259
 - staging parameters 376
 - Staging Parameters tab 234
- users
 - adding 17
- users, adding 16
- UTILEXCL

- customizations for the CICS domain
 - category 332
- customizations for the DB2 domain
 - category 332

V

- variables
 - specifying ranking for 150
- viewing
 - reports 10
 - template tables 77
- VMware
 - converting to local time 126
 - Guest and Host Disk, Network, and System domain categories 497
- VMware adapter
 - duplicate-data checking for 414
 - limiting data in staged tables 319
- VMware Data Acquisition
 - staging parameters 379
- VMware lookup table
 - populating 322
 - working with 321
- VMware vCenter
 - staging parameters 379
- VMware vCenter Data Acquisition
 - adapter 351
- VMware vCenter Data Acquisition job
 - running 353
 - working with 352

W

- Web Log
 - accessing raw data 315
 - all domain categories 500
 - performance metrics 426
 - staging parameters 380
- Web-based reporting 10
- Windows system data sources 70, 292
- wizards 43
 - supplied by SAS Data Integration Studio 45
 - supplied by SAS IT Resource Management 43
- workload profiling and characterization
 - metrics 419, 427

Z

- z/OS
 - additional requirements 31
- z/OS data sources 68, 291