

SAS[®] 9.3 Intelligence Platform Overview



The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2011. *SAS® 9.3 Intelligence Platform: Overview*. Cary, NC: SAS Institute Inc.

SAS® 9.3 Intelligence Platform: Overview

Copyright © 2011, 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.

The scanning, uploading, and distribution of this book via the Internet or any other means without the permission of the publisher is illegal and punishable by law. Please purchase only authorized electronic editions and do not participate in or encourage electronic piracy of copyrighted materials. Your support of others' rights is appreciated.

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, July 2011

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>Recommended Reading</i> | v |
| Chapter 1 • Value of the SAS Intelligence Platform | 1 |
| About the SAS Business Analytics Framework and the SAS Intelligence Platform | 1 |
| Accessibility Features in the SAS Intelligence Platform Products | 2 |
| Components of the SAS Intelligence Platform | 2 |
| Strategic Benefits of the SAS Intelligence Platform | 7 |
| Chapter 2 • Architecture of the SAS Intelligence Platform | 9 |
| Architecture of the SAS Intelligence Platform | 9 |
| SAS Servers | 11 |
| Middle Tier | 13 |
| Clients | 14 |
| Chapter 3 • Data in the SAS Intelligence Platform | 15 |
| Overview of Data Storage Options | 15 |
| Default SAS Storage | 16 |
| Third-Party Data Storage | 16 |
| Parallel Storage | 17 |
| Multidimensional Databases (Cubes) | 18 |
| How Data Sources Are Managed in the Metadata | 18 |
| Chapter 4 • Servers in the SAS Intelligence Platform | 21 |
| Overview of SAS Servers | 21 |
| SAS Metadata Server | 23 |
| Server Objects, Application Servers, and Logical Servers | 25 |
| Load Balancing for SAS Workspace Servers, SAS Stored Process Servers, and SAS OLAP Servers | 27 |
| Workspace Pooling for SAS Workspace Servers | 27 |
| Logging and Monitoring for SAS Servers | 28 |
| SAS Grid Computing | 29 |
| Chapter 5 • Middle-Tier Components of the SAS Intelligence Platform | 31 |
| Overview of Middle-Tier Components | 31 |
| Third-Party Software Components | 32 |
| SAS Web Infrastructure Platform | 33 |
| SAS Content Server | 34 |
| SAS Foundation Services | 35 |
| Chapter 6 • Clients in the SAS Intelligence Platform | 37 |
| Overview of Clients | 37 |
| SAS Add-In for Microsoft Office | 38 |
| SAS BI Dashboard | 39 |
| SAS BI Portlets | 39 |
| SAS Data Integration Studio | 40 |
| SAS Enterprise Guide | 40 |
| SAS Enterprise Miner | 41 |
| SAS Forecast Studio | 41 |
| SAS Information Delivery Portal | 42 |

| | |
|-----------------------------------------------------------|-----------|
| SAS Information Map Studio | 42 |
| JMP | 42 |
| SAS Management Console | 43 |
| SAS Model Manager | 44 |
| SAS OLAP Cube Studio | 44 |
| SAS Web Parts for Microsoft SharePoint | 45 |
| SAS Web Report Studio | 45 |
| SAS Workflow Studio | 46 |
| SAS Help Viewer for the Web | 46 |
| Chapter 7 • Security Overview | 47 |
| Overview of SAS Intelligence Platform Security | 47 |
| Authorization and Permissions Overview | 48 |
| Roles and Capabilities Overview | 49 |
| Authentication and Identity Management Overview | 49 |
| Single Sign-On in the SAS Intelligence Platform | 50 |
| Overview of Initial Users | 51 |
| Encryption Overview | 51 |
| Security Reporting and Logging Overview | 52 |
| Glossary | 53 |
| Index | 63 |

Recommended Reading

Here is the recommended reading list for this title:

- *SAS Intelligence Platform: Installation and Configuration Guide*
- *SAS Intelligence Platform: Migration Guide*
- *SAS Intelligence Platform: System Administration Guide*
- *SAS Intelligence Platform: Security Administration Guide*
- *SAS Intelligence Platform: Application Server Administration Guide*
- *SAS Intelligence Platform: Desktop Application Administration Guide*
- *SAS Intelligence Platform: Data Administration Guide*
- *SAS Intelligence Platform: Middle-Tier Administration Guide*
- *SAS Intelligence Platform: Web Application Administration Guide*
- *Grid Computing in SAS*
- *SAS BI Dashboard: User's Guide*
- *SAS Data Integration Studio: User's Guide*
- *SAS Foundation Services: Administrator's Guide*
- *SAS High-Performance Forecasting: User's Guide*
- *SAS Information Map Studio: Getting Started with SAS Information Maps*
- *SAS Logging: Configuration and Programming Reference*
- *SAS OLAP Server: MDX Guide*
- *SAS OLAP Server: User's Guide*
- *SAS Providers for OLE DB: Cookbook*
- *SAS Scalable Performance Data Engine: Reference*
- *SAS Web Report Studio: User's Guide*
- *SAS/ACCESS for Relational Databases: Reference*
- *Scheduling in SAS*

For a complete list of SAS publications, go to support.sas.com/bookstore. If you have questions about which titles you need, please contact a SAS Publishing Sales Representative:

SAS Publishing Sales

vi *Recommended Reading*

SAS Campus Drive
Cary, NC 27513-2414
Phone: 1-800-727-3228
Fax: 1-919-677-8166
E-mail: sasbook@sas.com
Web address: support.sas.com/bookstore

Chapter 1

Value of the SAS Intelligence Platform

| | |
|-------------------------------------------------------------------------------------------|----------|
| About the SAS Business Analytics Framework and the SAS Intelligence Platform | 1 |
| Accessibility Features in the SAS Intelligence Platform Products | 2 |
| Components of the SAS Intelligence Platform | 2 |
| Components Overview | 2 |
| Data Management Overview | 3 |
| Business Intelligence | 5 |
| Analytics | 6 |
| Supporting Components | 7 |
| Strategic Benefits of the SAS Intelligence Platform | 7 |
| Multiple Capabilities Integrated into One Platform | 7 |
| Consistency of Data and Business Rules | 8 |
| Fast and Easy Reporting and Analysis | 8 |
| Analytics Available to All Users | 8 |

About the SAS Business Analytics Framework and the SAS Intelligence Platform

The SAS Business Analytics Framework encompasses the full range of business solutions, technologies, and services from SAS. Through this framework, organizations can address their most critical business issues and then add new functionality over time to enable continuous performance improvement. All of the functionality is available from one vendor and through one framework, thus reducing the total cost of ownership.

The technologies that provide the foundation for the SAS Business Analytics Framework, as well as for industry and line-of-business solutions offered by SAS, are delivered through the SAS Intelligence Platform. The SAS Intelligence Platform is a comprehensive, end-to-end infrastructure for creating, managing, and distributing enterprise intelligence. It includes tools and interfaces that enable you to do the following:

- extract data from a variety of operational data sources on multiple platforms, and build a data warehouse and data marts that integrate the extracted data
- store large volumes of data efficiently and in a variety of formats
- give business users at all levels the ability to explore data from the warehouse in a Web browser, perform simple query and reporting functions, and view up-to-date results of complex analyses

- use high-end analytic techniques to provide capabilities such as predictive and descriptive modeling, forecasting, optimization, simulation, and experimental design
- centrally control the accuracy and consistency of enterprise data

With the SAS Intelligence Platform, you can implement an end-to-end intelligence infrastructure using software that is delivered, tested, and integrated by SAS. Using the tools provided in the SAS Intelligence Platform, you can create applications that reflect your unique business requirements and domain knowledge.

Building on the technologies in the SAS Intelligence Platform, SAS offers solutions for industries such as financial services, life sciences, health care, retail, and manufacturing, as well as line-of-business solutions in areas such as customer management, enterprise risk management, and financial management. The solutions incorporate predictive analytics, industry and domain expertise, and specialized data structures.

A complete set of documentation for deploying and administering the SAS Intelligence Platform is available at <http://support.sas.com/93administration>. The site provides the primary administrative documentation for the following software offerings:

- SAS BI Server
- SAS Enterprise BI Server
- SAS Data Integration Server
- SAS Enterprise Data Integration Server

Administrators who manage SAS products that use the metadata server should use the documentation at this site. Supplemental guides provide additional details for other products and domains.

Accessibility Features in the SAS Intelligence Platform Products

For information about accessibility for any of the products mentioned in this book, see the documentation for that product. If you have questions or concerns about the accessibility of SAS products, send e-mail to accessibility@sas.com.

Components of the SAS Intelligence Platform

Components Overview

The SAS Intelligence Platform includes components in the following categories:

Data Management

The data management components enable you to consolidate and manage enterprise data from a variety of source systems, applications, and technologies. Components are provided to help you cleanse, migrate, synchronize, replicate, and promote your data. In addition, SAS offers data storage options that are optimized for analytical processing, enabling you to quickly analyze and report on large volumes of data. Metadata for all of your intelligence resources is stored centrally and controlled through a single management interface.

Business Intelligence

The business intelligence components enable users with various needs and skill levels to create, produce, and share their own reports and analyses. Through easy-to-use interfaces, users can obtain their own answers to business questions. Meanwhile, the information technology staff retains control over the quality and consistency of the data.

Analytics

SAS offers the richest and widest portfolio of analytic products in the software industry. The portfolio includes products for predictive and descriptive modeling, data mining, text analytics, forecasting, optimization, simulation, data visualization, model management, and experimental design. You can use any combination of these tools with the SAS Intelligence Platform to add extraordinary precision and insight to your reports and analyses.

The following sections describe in more detail the data management, business intelligence, and analytics components, as well as key supporting components.

Data Management Overview

Data Management

The software tools in the data management category enable you to consolidate and manage enterprise data from a variety of source systems, applications, and technologies. SAS provides access engines and interfaces to a wide variety of data sources, including the following:

- delimited files, SAS data sets, and relational database management system (RDMS) tables
- application data from enterprise resource planning (ERP) and customer relationship management (CRM) systems
- message queuing platforms
- Web services
- unstructured and semi-structured data

Data storage options include simple relational databases, a threaded multidimensional database that supports online analytical processing (OLAP), and relational storage with a threaded multiple input/output (I/O) subsystem for intensive use by focused applications.

Each of the data management solutions is described briefly in the following sections.

SAS Data Integration Server and SAS Data Integration Studio

SAS Data Integration Studio, which is part of the SAS Data Integration Server offering, is a visual design tool that enables you to consolidate and manage enterprise data from a variety of source systems, applications, and technologies. The software enables you to create jobs and process flows that extract, transform, and load data for use in data warehouses and data marts. You can also create processes that cleanse, migrate, synchronize, replicate, and promote data for applications and business services.

For more information, see [“Clients in the SAS Intelligence Platform” on page 37](#) .

DataFlux Data Management Platform

The DataFlux Data Management Platform works with SAS Data Integration Studio and enables you to discover, design, deploy, and maintain data across your enterprise in a

centralized way. The DataFlux Data Management Platform combines data quality, data integration, and master data management under a unified user interface.

SAS Data Surveyor

The SAS Data Surveyor applications enable you to build SAS Data Integration Studio jobs to read data directly from these ERP vendors: SAP, Oracle, Salesforce.com, and Siebel.

SAS Data Quality Server

The SAS Data Quality Server works with software from SAS and DataFlux (a SAS company) to analyze, cleanse, transform, and standardize your data. The language elements that make up the SAS Data Quality Server software form the basis of the data quality transformations in SAS Data Integration Studio.

Data Storage Options

The data storage options that can be used with the SAS Intelligence Platform include SAS data tables, parallel storage, multidimensional databases, and third-party databases. These storage options can be used alone or in any combination. Metadata for your intelligence resources is stored centrally in the SAS Metadata Repository for use by all components of the intelligence platform.

Relational Storage: SAS Data Sets

You can use SAS data sets, the default SAS storage format, to store data of any granularity. The data values in a SAS data set are organized as a table of observations (rows) and variables (columns). A SAS data set also contains descriptor information such as the data types and lengths of the columns, as well as which SAS engine was used to create the data.

Access to Third-Party Databases: SAS/ACCESS

SAS/ACCESS provides interfaces to a wide range of relational, hierarchical, and network model databases. Examples include DB2, Oracle, SQL Server, Teradata, IBM Information Management System (IMS), and Computer Associates Integrated Database Management System (CA-IDMS). With SAS/ACCESS, SAS Data Integration Studio and other SAS applications can read, write, and update data regardless of which database and platform the data is stored on. SAS/ACCESS interfaces provide fast, efficient data loading and enable SAS applications to work directly from your data sources without making a copy.

Several SAS/ACCESS engines support bulk load of data files and threaded reads. Threaded reads let you read blocks of data on multiple threads instead of a record at a time. Several engines can use multiple threads to the parallel database management system (DBMS) server to access DBMS data. Both features significantly improve performance so that you can read and load data more rapidly without changing your original data.

High-Performance Computing: SAS In-Database

To support high-performance computing for complex, high-volume analytics, SAS In-Database enables certain data management, analytic, and reporting tasks to be performed inside the database. In-database technology minimizes the movement of data across the network, while enabling more sophisticated queries and producing results more quickly. This technology is available for several types of databases.

Multidimensional Storage: SAS OLAP Server

The SAS OLAP Server provides dedicated storage for data that has been summarized along multiple business dimensions. The server uses a threaded, scalable, and open technology and is especially designed for fast-turnaround processing and reporting.

A simplified ETL process enables you to build consistent OLAP cubes from disparate systems. A threaded query engine and parallel storage enable data to be spread across multiple-disk systems. Support is provided for multidimensional (MOLAP) and hybrid (HOLAP) data stores, as well as for open industry standards.

Parallel Storage: SAS Scalable Performance Data Engine and SAS Scalable Performance Data Server

The SAS SPD Engine and SAS SPD Server provide a high-speed data storage alternative for processing very large SAS data sets. They read and write tables that contain millions of observations, including tables that exceed the 2-GB size limit imposed by some operating systems. In addition, they provide the rapid data access that is needed to support intensive processing by SAS analytic software and procedures.

These facilities work by organizing data into a streamlined file format and then using threads to read blocks of data very rapidly and in parallel. The software tasks are performed in conjunction with an operating system that enables threads to execute on any of the CPUs that are available on a machine.

The SAS SPD Engine, which is included with Base SAS software, is a single-user data storage solution. The SAS SPD Server, which is available as a separate product, is a multi-user solution that includes a comprehensive security infrastructure, backup and restore utilities, and sophisticated administrative and tuning options.

Business Intelligence

The software tools in the business intelligence category address two main functional areas: information design, and self-service reporting and analysis.

The information design tools enable business analysts and information architects to organize data in ways that are meaningful to business users, while shielding the end users from the complexities of underlying data structures. These tools include the following products:

- SAS Information Map Studio enables analysts and information architects to create and manage information maps that contain business metadata about your data.
- SAS OLAP Cube Studio enables information architects to create cube definitions that organize summary data along multiple business dimensions.

The self-service reporting and analysis tools enable business users to query, view, and explore centrally stored information. Users can create their own reports, graphs, and analyses in the desired format and level of detail. In addition, they can find, view, and share previously created reports and analyses. The tools feature intuitive interfaces that enable business users to perform these tasks with minimal training and without the involvement of information technology staff.

The self-service reporting and analysis tools include the following products:

- SAS Web Report Studio is a Web-based query and reporting tool that enables users at any skill level to create, view, and organize reports.
- SAS Information Delivery Portal provides a Web-based, personalized workplace to help decision makers easily find the information that they need.
- SAS BI Portlets includes portlets, such as the SAS Stored Process Portlet and the SAS Report Portlet, that add value to the SAS Information Delivery Portal.
- SAS BI Dashboard enables SAS Information Delivery Portal users to create, maintain, and view dashboards to monitor key performance indicators that convey how well an organization is performing.

- SAS Add-In for Microsoft Office enables users to access SAS functionality from within Microsoft Office products.
- SAS Enterprise Guide is a project-oriented Windows application that enables users to create processes that include complex computations, business logic, and algorithms.

As users create information maps, cubes, report definitions, portal content definitions, and stored processes, information about them is stored in the SAS Metadata Repository. Client applications and users can access these information assets on a need-to-know basis. Access is controlled through multilayered security that is enforced through the metadata.

For a description of each of the business intelligence tools, see [“Clients in the SAS Intelligence Platform” on page 37](#).

Analytics

SAS offers the richest and widest portfolio of analytic products in the software industry. The portfolio includes products for predictive and descriptive modeling, data mining, text analytics, forecasting, optimization, simulation, and experimental design. You can use any combination of these tools with the SAS Intelligence Platform to add precision and insight to your reports and analyses.

SAS software provides the following types of analytical capabilities:

- predictive analytics and data mining, to build descriptive and predictive models and deploy the results throughout the enterprise
- text analytics, to maximize the value of unstructured data assets
- dynamic data visualization, to enhance the effectiveness of analytics
- forecasting, to analyze and predict outcomes based on historical patterns
- model management, to streamline the process of creating, managing, and deploying analytic models
- operations research, to apply techniques such as optimization, scheduling, and simulation to achieve the best result
- quality improvement, to identify, monitor, and measure quality processes over time
- statistical data analysis, to drive fact-based decisions

The following are examples of analytic products:

- SAS Enterprise Miner enables analysts to create and manage data mining process flows. These flows include steps to examine, transform, and process data to create models that predict complex behaviors of economic interest. The SAS Intelligence Platform enables SAS Enterprise Miner users to centrally store and share the metadata for models and projects. In addition, SAS Data Integration Studio provides the ability to schedule data mining jobs.
- SAS Forecast Server enables organizations to plan more effectively for the future by generating large quantities of high-quality forecasts quickly and automatically. This solution includes the SAS High-Performance Forecasting engine, which selects the time series models, business drivers, and events that best explain your historical data, optimizes all model parameters, and generates high-quality forecasts. SAS Forecast Studio provides a graphical interface to these high-performance forecasting procedures.

- SAS Model Manager supports the deployment of analytical models into your operational environments. It enables registration, modification, tracking, scoring, and reporting on analytical models that have been developed for BI and operational applications.
- JMP is interactive, exploratory data analysis and modeling software for the desktop. JMP makes data analysis—and the resulting discoveries—visual and helps to communicate those discoveries to others. JMP presents results both graphically and numerically. By linking graphs to each other and to the data, JMP makes it easier to see the trends, outliers, and other patterns that are hidden in your data.

Supporting Components

SAS Metadata Repository

Your information assets are managed in a common metadata layer called the SAS Metadata Repository.

This repository stores logical data representations of items such as libraries, tables, information maps, and cubes, thus ensuring central control over the quality and consistency of data definitions and business rules. The repository also stores information about system resources such as servers, the users who access data and metadata, and the rules that govern who can access what.

All of the data management and business intelligence tools read and use metadata from the repository and create new metadata as needed.

SAS Management Console

SAS Management Console provides a single interface through which system administrators can manage and monitor SAS servers, explore and manage metadata repositories, manage user and group accounts, and administer security.

Scheduling in SAS

Platform Suite for SAS is an optional product that provides enterprise-level scheduling capabilities in a single-server environment. Platform Suite for SAS is also included as part of the SAS Grid Manager product to enable distributed enterprise scheduling, workload balancing, and parallelized workload balancing. The components of Platform Suite for SAS include Process Manager, Load Sharing Facility (LSF), and Grid Management Services.

As an alternative, operating system services can be used to provide a basic level of scheduling for SAS jobs, and SAS in-process scheduling enables you to schedule jobs from certain Web-based SAS applications.

Strategic Benefits of the SAS Intelligence Platform

Multiple Capabilities Integrated into One Platform

The SAS Intelligence Platform combines advanced SAS analytics, high-speed processing of large amounts of data, and easy-to-use query and reporting tools. The result is accurate, reliable, and fast information with which to make decisions.

You can build data warehouses, perform data mining, enable users to query data and produce reports from a Web browser, and give users easy access to SAS processes that perform robust analytics.

The SAS Intelligence Platform provides all of this functionality in one centrally managed suite of products that are designed to work together seamlessly. This integration reduces the administration, management, and deployment costs that would be associated with providing multiple technologies to meet the needs of different users.

Consistency of Data and Business Rules

The SAS Intelligence Platform makes use of your organization's existing data assets, enabling you to integrate data from multiple database platforms and ERPs. Tools are provided to help ensure the reliability, consistency, and standardization of this data.

Users can choose from multiple tools with which to perform queries and produce reports. Since all of the tools access data through the same metadata representations, users throughout your enterprise receive consistent data. As a result, they can make decisions based on a common version of the truth.

Similarly, business logic, complex computations, and analytic algorithms can be developed once and stored centrally in SAS processes for all users to access. These processes, as well as the information in the SAS Metadata Repository, are controlled through multi-level security.

Fast and Easy Reporting and Analysis

The SAS Intelligence Platform's self-service reporting and analysis tools enable users across the enterprise to access and query data from virtually any data source. Any number of users can use wizards to create reports in the needed time frames, without waiting for support from information technology professionals. Through Web-based interfaces, users can explore large volumes of multidimensional data quickly and interactively, from multiple perspectives and at multiple levels of detail.

The reporting and analysis tools hide complex data structures, so that average business users can perform queries without having to learn new skills. The intelligence storage options are optimized for analytical processing, enabling the reporting tools to quickly retrieve large volumes of data.

As a result of these reporting and analysis capabilities, everyone spends less time looking for answers and more time driving strategic decisions.

Analytics Available to All Users

SAS is the market leader in analytics. With the SAS Intelligence Platform, you can make the full breadth of SAS analytics available to users throughout the enterprise.

SAS analytics include algorithms for functions such as predictive and descriptive modeling, forecasting, optimization, simulation, and experimental design. You can now incorporate these capabilities into self-service reports and analyses, so that decision makers throughout your enterprise can benefit from the accuracy and precision of high-end analytics.

Chapter 2

Architecture of the SAS Intelligence Platform

| | |
|----------------------------------------------------------------|-----------|
| Architecture of the SAS Intelligence Platform | 9 |
| Architecture Overview | 9 |
| Data Sources | 10 |
| SAS Servers | 11 |
| SAS Servers Overview | 11 |
| SAS Metadata Server | 12 |
| SAS OLAP Server | 12 |
| SAS Workspace Server | 12 |
| SAS Pooled Workspace Server | 13 |
| SAS Stored Process Server | 13 |
| SAS Object Spawner | 13 |
| Middle Tier | 13 |
| Clients | 14 |

Architecture of the SAS Intelligence Platform

Architecture Overview

The SAS Intelligence Platform architecture is designed to efficiently access large amounts of data, while simultaneously providing timely intelligence to a large number of users. The platform uses an n-tier architecture that enables you to distribute functionality across computer resources, so that each type of work is performed by the resources that are best suited to the job.

You can easily scale the architecture to meet the demands of your workload. For a large company, the tiers can be installed across a multitude of machines with different operating systems. For prototyping, demonstrations, or very small enterprises, all of the tiers can be installed on a single machine.

The architecture consists of the following four tiers:

Data sources

Data sources store your enterprise data. All of your existing data assets can be used, whether your data is stored in third-party database management systems, SAS tables, or enterprise resource planning (ERP) system tables.

SAS servers

SAS servers perform SAS processing on your enterprise data. Several types of SAS servers are available to handle different workload types and processing intensities.

The software distributes processing loads among server resources so that multiple client requests for information can be met without delay.

Middle tier

The middle tier enables users to access intelligence data and functionality via a Web browser. This tier provides Web-based interfaces for report creation and information distribution, while passing analysis and processing requests to the SAS servers.

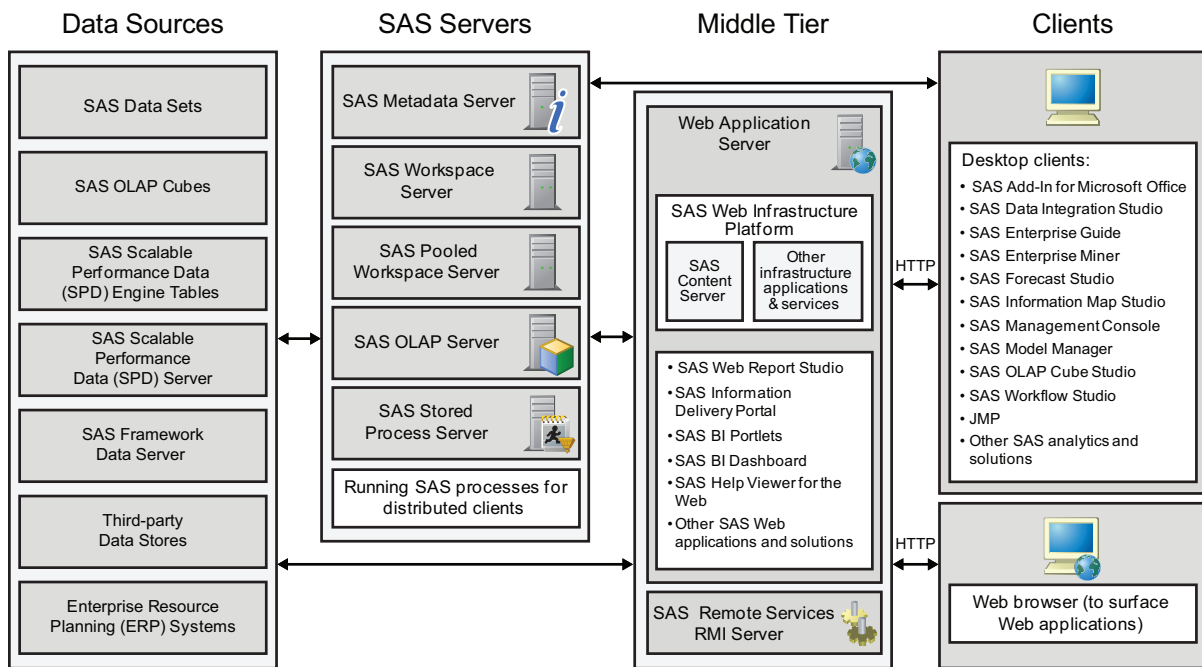
Clients

The client tier provides users with desktop access to intelligence data and functionality through easy-to-use interfaces. For most information consumers, reporting and analysis tasks can be performed with just a Web browser. For more advanced design and analysis tasks, SAS client software is installed on users' desktops.

Note: The four tiers listed above represent categories of software that perform similar types of computing tasks and require similar types of resources. The tiers do not necessarily represent separate computers or groups of computers.

The following diagram shows how the tiers interact, and the sections that follow describe each tier in more detail.

Figure 2.1 Architecture of the SAS Intelligence Platform



Data Sources

The SAS Intelligence Platform includes the following options for data storage:

- SAS data sets, which are analogous to relational database tables.
- SAS SPD Engine tables, which can be read or written by multiple threads.
- SAS SPD Server, which is available as a separate product.
- SAS OLAP cubes.

- The SAS Framework Data Server, which is the default location for middle-tier data such as alerts, comments, and workflows, as well as data for the SAS Content Server and SAS Service Parts Optimization. The server is provided as an alternative to using a third-party DBMS. (The server cannot be used as a general-purpose data store.)

In addition, SAS provides products that enable you to access data in your existing third-party data stores and ERP systems. The SAS/ACCESS interfaces provide direct access to data stores such as the following:

- Oracle
- DB2
- Teradata
- ODBC
- Microsoft SQL Server
- Microsoft Excel spreadsheets
- Sybase

For a complete list of supported data stores, go to <http://support.sas.com/software/products/access/>.

The following SAS Data Surveyor products provide direct access to ERP systems:

- SAS Data Surveyor for SAP Plug-in for SAS Data Integration Studio enables you to consolidate and manage enterprise data from various SAP R/3 data sources.
- SAS Data Surveyor for Clickstream Data is a plug-in to SAS Data Integration Studio. This plug-in enables you to create jobs that extract and transform clickstream data from Web logs, and then load the resulting data into a SAS table. Other applications, such as SAS Web Analytics, can then take the refined clickstream data and analyze it.
- SAS has partnered with Composite Software to provide customers with access to ERP and customer relationship management (CRM) data. SAS data surveyors for Oracle Applications, Salesforce.com, and Siebel are available through this partnership.

For more information about data sources, see “Data in the SAS Intelligence Platform” on page 15 .

SAS Servers

SAS Servers Overview

The SAS servers execute SAS analytical and reporting processes for distributed clients. These servers are typically accessed either by desktop clients or by Web applications that are running in the middle tier.

Note: In the SAS Intelligence Platform, the term server refers to a process or processes that wait for and fulfill requests from client programs for data or services. The term server does not necessarily refer to a specific computer, since a single computer can host one or more servers of various types.

The SAS servers use the SAS Integrated Object Model (IOM), which is a set of distributed object interfaces that make SAS software features available to client

applications when SAS is executed on a server. Each server uses a different set of IOM interfaces and has a different purpose.

The principal servers in the SAS Intelligence Platform include the SAS Metadata Server, the SAS OLAP Server, the SAS Workspace Server, the SAS Pooled Workspace Server, the SAS Stored Process Server, and the SAS Framework Data Server. The SAS object spawner process runs on server machines and handles requests for SAS Workspace Servers, SAS Pooled Workspace Servers, and SAS Stored Process Servers.

SAS Metadata Server

The SAS Metadata Server controls access to a central repository of metadata that is shared by all of the SAS applications in the deployment. The SAS Metadata Server enables centralized control so that all users access consistent and accurate data. The metadata repository stores information about the following:

- the enterprise data sources and data structures that are accessed by SAS intelligence applications.
- the content that is created and used by SAS applications. This content includes information maps, OLAP cubes, report definitions, stored process definitions, and portal content definitions.
- the SAS and third-party servers that participate in the system.
- users and groups of users who are allowed to use the system. Users can be authenticated by the metadata server or by external systems such as the host environment, the Web realm, and third-party databases.
- the levels of access that users and groups have to resources. This metadata-based authorization layer supplements protections from the host environment and other systems.

The SAS Intelligence Platform provides a central management tool—SAS Management Console—that you use to manage the metadata server and the metadata repository.

SAS OLAP Server

The SAS OLAP Server is a multidimensional data server that delivers pre-summarized cubes of data to business intelligence applications. The data is queried using the MDX (multidimensional expression) language.

This server is designed to reduce the load on traditional back-end storage systems by quickly delivering summarized views, irrespective of the amount of data that underlies the summaries.

SAS Workspace Server

The SAS Workspace Server enables client applications to submit SAS code to a SAS session using an application programming interface (API). For example, when you use SAS Data Integration Studio to submit an ETL job for processing, the application generates the SAS code necessary to perform the processing and submits it to a workspace server. You can run as many instances of workspace servers as are needed to support your workload.

SAS Pooled Workspace Server

The SAS Pooled Workspace Server is a workspace server that uses server-side pooling. This configuration maintains a collection of reusable workspace server processes for clients, thus avoiding the overhead associated with creating a new process for each connection. Clients such as SAS Information Map Studio, SAS Web Report Studio, and the SAS Information Delivery Portal can use pooled workspace servers to query relational data.

SAS Stored Process Server

The SAS Stored Process Server executes and delivers results from SAS Stored Processes in a multi-client environment. A SAS Stored Process is a SAS program that is stored centrally and that can be executed by users and client programs on demand.

You can run as many instances of stored process servers as are needed to support your workload.

SAS Object Spawner

The SAS object spawner is a process that runs on workspace server, pooled workspace server, and stored process server host machines. It listens for requests for these servers, authenticates the requesting clients, and launches server processes as needed. In a pooled workspace server configuration, the object spawner maintains a collection of re-usable workspace server processes that are available for clients. If server load balancing is configured, the object spawner balances workloads between server processes. The object spawner connects to the metadata server to obtain information about the servers that it manages.

Middle Tier

The middle tier of the SAS Intelligence Platform provides an environment in which the business intelligence Web applications, such as SAS Web Report Studio and the SAS Information Delivery Portal, can execute. These products run in a Web application server and communicate with the user by sending data to and receiving data from the user's Web browser. The middle tier applications rely on servers on the SAS server tier to perform SAS processing, including data query and analysis.

The middle tier includes the following third-party software and SAS software elements:

- a Web application server
- the Java Development Kit (JDK)
- SAS Web applications, which can include SAS Web Report Studio, the SAS Information Delivery Portal, SAS BI Portlets, the SAS BI Dashboard, SAS Help Viewer for the Web, and other SAS products and solutions
- the SAS Web Infrastructure Platform, which includes the SAS Content Server and other infrastructure applications and services
- a Java remote method invocation (RMI) server, which enables access to SAS Foundation Services and associated extension services

For more information about the middle tier, see [“Middle-Tier Components of the SAS Intelligence Platform”](#) on page 31.

Clients

The clients in the SAS Intelligence Platform provide Web-based and desktop user interfaces to content and applications. SAS clients provide access to content, appropriate query and reporting interfaces, and business intelligence functionality for all of the information consumers in your enterprise, from the CEO to business analysts to customer service agents.

The following clients run on Windows desktops. Some of these clients are native Windows applications and others are Java applications.

- SAS Add-In for Microsoft Office
- SAS Data Integration Studio
- SAS Enterprise Guide
- SAS Enterprise Miner
- SAS Forecast Studio
- SAS Information Map Studio
- JMP (also available on Macintosh and Linux)
- SAS Management Console
- SAS Model Manager
- SAS OLAP Cube Studio
- SAS Workflow Studio

SAS Management Console is supported on all platforms except z/OS.

The following products require only a Web browser to be installed on each client machine, with the addition of Adobe Flash Player for SAS BI Dashboard:

- SAS Information Delivery Portal
- SAS BI Dashboard
- SAS Web Report Studio
- SAS Help Viewer for the Web

For more information about the clients, see [“Clients in the SAS Intelligence Platform”](#) on page 37.

Chapter 3

Data in the SAS Intelligence Platform

| | |
|---------------------------------------------------------------|-----------|
| Overview of Data Storage Options | 15 |
| Default SAS Storage | 16 |
| Third-Party Data Storage | 16 |
| Parallel Storage | 17 |
| Parallel Storage Overview | 17 |
| Options for Implementing Parallel Storage | 17 |
| How Parallel Storage Works | 17 |
| Multidimensional Databases (Cubes) | 18 |
| How Data Sources Are Managed in the Metadata | 18 |

Overview of Data Storage Options

In a SAS Intelligence Platform deployment, you can use one or more of these data storage options:

- [default SAS storage](#) in the form of SAS tables
- [third-party relational, hierarchical, and network model database tables](#)
- [parallel storage](#) from the SAS Scalable Performance Data Engine (SPD Engine) and the SAS Scalable Performance Data Server (SPD Server)
- [multidimensional databases \(cubes\)](#)

All four data sources provide input to reporting applications. The first three sources are also used as input for these data structures:

- cubes, which are created with either SAS Data Integration Studio or SAS OLAP Cube Studio
- data marts and data warehouses, which are created with SAS Data Integration Studio

You can use these storage options in any combination to meet your unique business requirements. The following sections describe each storage option in more detail. Central [management of data sources](#) through the SAS metadata repository is also discussed.

Default SAS Storage

You can use SAS data sets (tables), the default SAS storage format, to store data of any granularity.

A SAS table is a file that SAS software creates and processes. Each SAS table is a member of a SAS library. A SAS library is a collection of one or more SAS files that are recognized by SAS software and that are referenced and stored as a unit.

Each SAS table contains the following:

- data values that are organized as a table of observations (rows) and variables (columns) that can be processed by SAS software
- descriptor information such as data types, column lengths, and the SAS engine that was used to create the data

For shared access to SAS tables, you can use SAS/SHARE software, which provides concurrent Update access to SAS files for multiple users.

Third-Party Data Storage

Data can be stored in a wide range of third-party databases, including the following:

- relational databases such as Oracle, Sybase, DB2, SQL Server, and Teradata
- hierarchical databases such as IBM Information Management System (IMS)
- Computer Associates Integrated Database Management System (CA-IDMS), which is a network model database system

SAS/ACCESS interfaces provide fast, efficient loading of data to and from these facilities. With these interfaces, SAS software can work directly from the data sources without making a copy. Several of the SAS/ACCESS engines use an input/output (I/O) subsystem that enables applications to read entire blocks of data instead of reading just one record at a time. This feature reduces I/O bottlenecks and enables procedures to read data as fast as they can process it. The SAS/ACCESS engines for Oracle, Sybase, DB2 (on UNIX and PC), ODBC, SQL Server, and Teradata support this functionality. These engines, as well as the DB2 engine on z/OS, can also access database management system (DBMS) data in parallel by using multiple threads to the parallel DBMS server. Coupling the threaded SAS procedures with these SAS/ACCESS engines provides even greater gains in performance.

SAS In-Database enables high-performance computing for complex, high-volume analytics. This technology enables certain Base SAS and SAS/STAT procedures to run aggregations and analytics inside the database. In-database technology minimizes the movement of data across the network, while enabling more sophisticated queries and producing results more quickly. This technology is available through the use of the SAS/ACCESS, SAS Scoring Accelerator, and SAS Analytics Accelerator products. The supported databases include the following:

- Aster *n*Cluster
- DB2 (UNIX only)
- Greenplum

- Netezza
- Oracle
- Teradata

Parallel Storage

Parallel Storage Overview

The SAS Scalable Performance Data Engine (SPD Engine) and the SAS Scalable Performance Data Server (SPD Server) are designed for high-performance data delivery. They enable rapid access to SAS data for intensive processing by the application.

Although the Base SAS engine is sufficient for most tables that do not span volumes, the SAS SPD Engine and SAS SPD Server are high-speed alternatives for processing very large tables. They read and write tables that contain millions of observations, including tables that expand beyond the 2-GB size limit imposed by some operating systems. In addition, they support SAS analytic software and procedures that require fast processing of tables.

Options for Implementing Parallel Storage

Two options are available for implementing parallel storage:

- The SAS SPD Engine is included with Base SAS software. It is a single-user data storage solution that shares the high-performance parallel processing and parallel I/O capabilities of SAS SPD Server, but lacks the additional complexity of a multi-user server.

The SPD Engine runs in UNIX and Windows operating environments as well as some z/OS operating environments.

- The SAS SPD Server is available as a separate product. It is a multi-user parallel-processing data server with a comprehensive security infrastructure, backup and restore utilities, and sophisticated administrative and tuning options.

The SAS SPD Server runs in Windows and UNIX operating environments.

How Parallel Storage Works

The SAS SPD Engine and SAS SPD Server deliver data to applications rapidly by organizing large SAS data sets into a streamlined file format. The file format enables multiple CPUs and I/O channels to perform parallel input/output (I/O) functions on the data. Parallel I/O takes advantage of multiple CPUs and multiple controllers, with multiple disks per controller, to read or write data in independent threads.

One way to take advantage of the features of the SAS SPD Engine is through a hardware and software architecture known as symmetric multiprocessing (SMP). An SMP machine has multiple CPUs and an operating system that supports threads. These machines are usually configured with multiple controllers and multiple disk drives per controller.

When the SAS SPD Engine reads a data file, it launches one or more threads for each of the CPUs in the SMP machine. These threads then read data in parallel from multiple

disk drives, driven by one or more controllers per CPU. The SAS SPD Engine running on an SMP machine provides the capability to read and deliver much more data to an application in a given elapsed time.

For more information, see the SAS SPD Server documentation, which is available at <http://support.sas.com/documentation/onlinedoc/spds/>.

Multidimensional Databases (Cubes)

Multidimensional databases (cubes) are another storage option provided by the SAS Intelligence Platform. Cubes provide business users with multiple views of their data through drill-down capabilities.

Cubes are derived from source data such as SAS tables, SAS SPD Engine tables, and SAS/ACCESS database tables. To create cube definitions, and to build cubes based on these definitions, you can use the Cube Designer wizard, which is available from SAS Data Integration Studio and SAS OLAP Cube Studio.

Cubes are managed by the SAS OLAP Server, which is a multi-user, scalable server designed to store and access large volumes of data while maintaining system performance.

The SAS OLAP Server uses a SAS engine that organizes data into a streamlined file format. This file format enables the engine to rapidly deliver data to client applications. The engine also reads and writes partitioned tables, which enables it to use multiple CPUs to perform parallel I/O functions. The threaded model enables the SAS OLAP Server to create and query aggregations in parallel for fastest performance.

SAS business intelligence applications perform queries against the cubes by using the multidimensional expression (MDX) query language. Cubes can be accessed by client applications that are connected to the SAS OLAP Server with the following tools:

- the SQL pass-through facility for OLAP, which is designed to process MDX queries within the PROC SQL environment
- open access technologies such as OLE DB for OLAP and ADO MD

How Data Sources Are Managed in the Metadata

All of the data sources that are used in your deployment of the SAS Intelligence Platform are centrally controlled through metadata that is stored in the SAS Metadata Repository. In the metadata repository, you can create the following types of metadata objects to control and manage your data:

- database servers, which provide relational database services to clients
- SAS Application Servers, which perform SAS processes on data
- cubes
- OLAP schemas, which specify which groups of cubes a given SAS OLAP Server can access
- dimensions and measures in a cube
- libraries, which are collections of one or more files that are recognized by SAS software and that are referenced and stored as a unit

- the data sources (for example, SAS tables) that are contained in a library
- the columns that are contained in a data source

A variety of methods are available to populate the metadata repository with these objects, including the following:

- The data source design applications, SAS Data Integration Studio and SAS OLAP Cube Studio, automate the creation of all of the necessary metadata about your data sources. As you use these products to define warehouses, data marts, and cubes, the appropriate metadata objects are automatically created and stored in the metadata repository.
- You can use the following features of SAS Management Console to define data source objects:
 - The New Server Wizard enables you to easily define the metadata for your database servers and SAS Application Servers.
 - The Data Library Manager enables you to define database schemas for a wide variety of schema types. You can also use this feature to define libraries if you are not using SAS Data Integration Studio to define them.
 - The Register Tables feature enables you to import table definitions from external sources if you are not using SAS Data Integration Studio to create them.
 - On the SAS Folders tab, you can set permissions that secure access to metadata folders and objects. Because all SAS Intelligence Platform applications use the metadata server when accessing resources, permissions that are enforced by the metadata server provide an effective level of protection. These permissions supplement protections from the host environment and other systems. Therefore, a user's ability to perform a particular action is determined not only by metadata-based access controls but also by external authorization mechanisms such as operating system permissions and database controls.
- You can use the metadata LIBNAME engine to enforce data-related Read, Write, Create, and Delete access controls that have been defined in metadata.

After your data sources have been defined in metadata, you can use SAS Information Map Studio to create SAS Information Maps, which are business metadata about your data. Information maps enable you to surface your data in business terms that typical business users understand, while storing key information that is needed to build appropriate queries.

For detailed information about administering data sources, see the *SAS Intelligence Platform: Data Administration Guide*.

Chapter 4

Servers in the SAS Intelligence Platform

| | |
|---------------------------------------------------------------------------------------------------------|-----------|
| Overview of SAS Servers | 21 |
| SAS Metadata Server | 23 |
| About the SAS Metadata Server | 23 |
| About the Metadata in the SAS Metadata Repository | 23 |
| How the Metadata Server Controls System Access | 24 |
| How Metadata Is Created and Administered | 24 |
| How Business Metadata Is Organized | 25 |
| Server Objects, Application Servers, and Logical Servers | 25 |
| About Server Objects and Server Groupings | 25 |
| Purpose of the Application Server Grouping | 26 |
| Purpose of the Logical Server Grouping | 27 |
| Load Balancing for SAS Workspace Servers, SAS Stored Process Servers, and SAS OLAP Servers | 27 |
| Workspace Pooling for SAS Workspace Servers | 27 |
| Logging and Monitoring for SAS Servers | 28 |
| SAS Grid Computing | 29 |

Overview of SAS Servers

The SAS Intelligence Platform provides access to SAS functionality through the following specialized servers:

Note: In the SAS Intelligence Platform, the term “server” refers to a program or programs that wait for and fulfill requests from client programs for data or services. The term server does not necessarily refer to a specific computer, since a single computer can host one or more servers of various types.

- the SAS Metadata Server, which writes metadata objects to, and reads metadata objects from, SAS Metadata Repositories. These metadata objects contain information about all of the components of your system, such as users, groups, data libraries, servers, and user-created products such as reports, cubes, and information maps.
- SAS Workspace Servers, which provide access to SAS software features such as the SAS language, SAS libraries, the server file system, results content, and formatting services.

A program called the SAS object spawner runs on a workspace server's host machine. The spawner listens for incoming client requests and launches server instances as needed.

- SAS Pooled Workspace Servers, which are configured to use server-side pooling. In this configuration, the SAS object spawner maintains a collection of workspace server processes that are available for clients. This server configuration is intended for use by query and reporting tools such as SAS Information Map Studio, SAS Web Report Studio, and the SAS Information Delivery Portal.
- SAS Stored Process Servers, which fulfill client requests to execute SAS Stored Processes. Stored processes are SAS programs that are stored on a server and can be executed as required by requesting applications. By default, three load balanced SAS Stored Process Servers are configured.

The SAS object spawner runs on a stored process server's host machine, listens for incoming client requests, and launches server instances as needed.

- SAS OLAP Servers, which provide access to cubes. Cubes are logical sets of data that are organized and structured in a hierarchical multidimensional arrangement. Cubes are queried by using the multidimensional expression (MDX) language.
- the SAS Framework Data Server, which is the default location for middle-tier data such as alerts, comments, and workflows, as well as data for the SAS Content Server and SAS Service Parts Optimization. The server is provided as an alternative to using a third-party DBMS. (The server cannot be used as a general-purpose data store.)
- SAS/CONNECT servers, which provide computing resources on remote machines where SAS Integration Technologies is not installed.
- batch servers, which enable you to execute code in batch mode. There are three types of batch servers: DATA step batch servers, Java batch servers, and generic batch servers. The DATA step server enables you to run SAS DATA steps and procedures in batch mode. The Java server enables you to schedule the execution of Java code, such as the code that creates a SAS Marketing Automation marketing campaign. The generic server supports the execution of any other type of code.
- the SAS Content Server, which is part of the middle tier architecture. This server stores digital content (such as documents, reports, and images) that is created and used by SAS Web applications. For more information, see [“SAS Content Server” on page 34](#).

Note: For accessing specialized data sources, the SAS Intelligence Platform can also include one or more data servers. These might include the SAS Scalable Performance Data (SPD) Server and third-party database management system (DBMS) products. The SAS OLAP Server also provides some data server functionality. For information about data servers, see [“Data in the SAS Intelligence Platform” on page 15](#).

The following sections describe these concepts:

- the central role of the [SAS Metadata Server](#) in the management of the SAS Intelligence Platform
- the organizational principles that are used to manage SAS server resources, including [server objects, logical servers, and application servers](#)
- the use of [load balancing](#) (for stored process servers and workspace servers) and [workspace pooling](#) (for workspace servers)
- the facilities and capabilities that are provided for [server logging and monitoring](#)

SAS Metadata Server

About the SAS Metadata Server

The SAS Metadata Server is a multi-user server that enables users to read metadata from and write metadata to one or more SAS Metadata Repositories. This server is a centralized resource for storing, managing, and delivering metadata for all SAS applications across the enterprise.

About the Metadata in the SAS Metadata Repository

Here are examples of the types of metadata objects that can be stored in the SAS Metadata Repository:

- users
- groups of users
- data libraries
- tables
- jobs
- cubes
- documents
- information maps
- reports
- stored processes
- dashboards
- analytical models
- SAS Workspace Servers
- SAS Pooled Workspace Servers
- SAS Stored Process Servers
- SAS OLAP Servers

A metadata object is a set of attributes that describe a resource. Here are some examples:

- When a user creates a report in SAS Web Report Studio, a metadata object is created to describe the new report.
- When a user creates a data warehouse in SAS Data Integration Studio, a metadata object is created to describe each table included in the warehouse.
- When a system administrator defines a new instance of a SAS server, a metadata object is created to describe the server.

The specific attributes that a metadata object includes vary depending on the resource that is being described. For example, a table object can include attributes for the table's name and description, path specification, host type, and associated SAS Application Server.

The SAS Metadata Server uses the SAS Open Metadata Architecture, which provide common metadata services to SAS and other applications. Third parties can access metadata in the SAS Metadata Server by using an application programming interface (API) that is supplied by SAS. SAS Metadata Bridges are available to support metadata exchange with a variety of sources, including the Common Warehouse Metadata Model (CWM).

How the Metadata Server Controls System Access

The SAS Metadata Server plays an important role in the security of the SAS Intelligence Platform. It controls system access in the following ways:

SAS identities

For accountability, we recommend that you create an individual SAS identity for each person who uses the SAS environment. These identities enable you to make access distinctions and audit individual actions in the metadata layer. The identities also provide personal folders for each user. The metadata server maintains its own copy of each user ID for the purpose of establishing a SAS identity.

access controls

You can define metadata-based access controls that supplement protections provided by the host environment and other systems. The metadata-based controls enable you to manage access to OLAP data, to relational data (depending on the method by which the data is accessed), and to almost any metadata object (for example, reports, data definitions, information maps, jobs, stored processes, and server definitions).

roles

You can assign users and groups to roles that determine whether they can use application features such as menu items, plug-ins, and buttons. Roles are designed to make application functionality available to the appropriate types of users. For example, role memberships determine whether a user can see the Server Manager plug-in (in SAS Management Console), compare data (in SAS Enterprise Guide), or directly open an information map (in SAS Web Report Studio). Applications that support roles include the SAS Add-In for Microsoft Office, SAS Enterprise Guide, SAS Forecast Studio, SAS Management Console, and SAS Web Report Studio.

For more information about security in the SAS Intelligence Platform, see [“Security Overview” on page 47](#).

How Metadata Is Created and Administered

Metadata can be loaded to the SAS Metadata Server in several ways:

- The configuration process for the SAS Intelligence Platform automatically creates and stores metadata objects for the resources, such as servers, that are part of your initial installation.
- SAS Metadata Bridges enable you to import metadata from a variety of sources, including market-leading design tool and repository vendors and the Common Warehouse Metamodel (CWM).
- When users create content such as reports, information maps, and data warehouses with the SAS Intelligence Platform applications, these applications create and store metadata objects describing the content.
- System administrators use the SAS Management Console client application to directly create metadata for system resources such as servers, users, and user groups.

System administrators also use SAS Management Console for general metadata administration tasks, including backing up the metadata server, creating new repositories, promoting metadata objects, and maintaining authorization information and access rules for all types of resources.

How Business Metadata Is Organized

A hierarchical folder structure is used to organize metadata for business intelligence content such as libraries, tables, jobs, cubes, information maps, and reports. The folder structure includes personal folders for individual users and a folder for shared data. Within this overall structure, you can create a customized folder structure that meets the information management, data sharing, and security requirements of your organization.

The Folders tab of SAS Management Console displays all SAS folders that the user has permission to view. Most other client applications display SAS folders only if they contain content that is relevant to the application, subject to the user's permissions. Import and export wizards in SAS Management Console, SAS Data Integration Studio, and SAS OLAP Cube Studio enable you to easily move or promote content from one part of the SAS Folders tree to another, or from a test environment to a production environment.

For more information about the SAS Metadata Server, see the *SAS Intelligence Platform: System Administration Guide*.

Server Objects, Application Servers, and Logical Servers

About Server Objects and Server Groupings

In the SAS Metadata Repository, each server process that executes SAS code is represented by a server object. In the metadata, the attributes for each server object contain information such as the following:

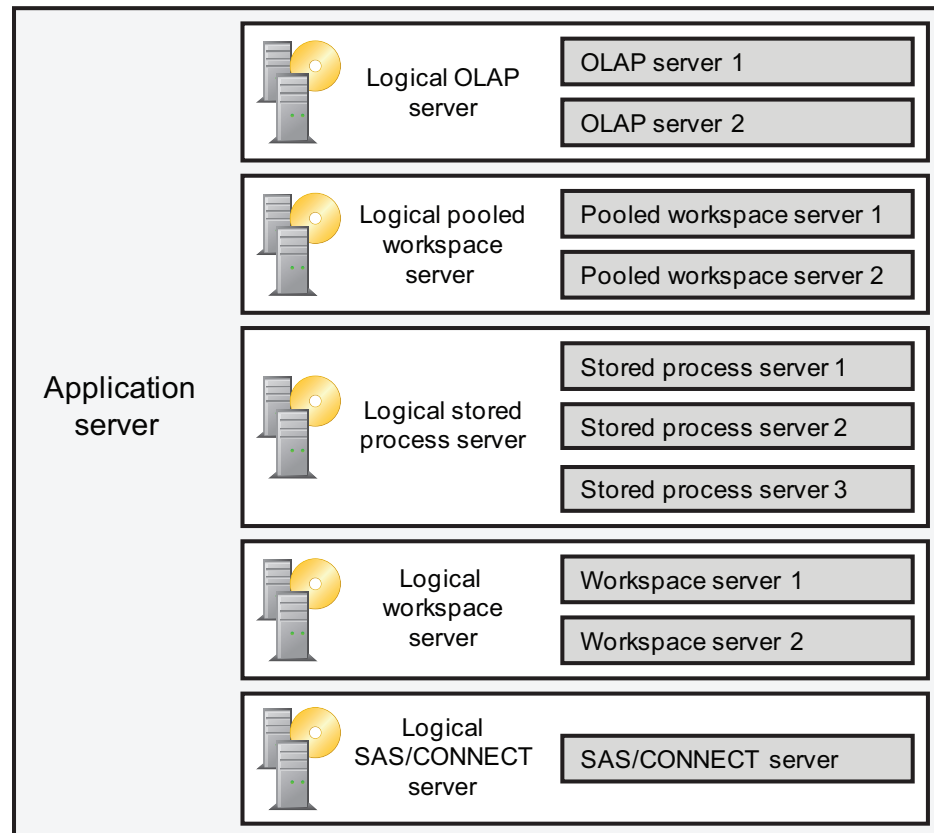
- the name of the machine that is hosting the server
- the TCP/IP port or ports on which the server listens for requests
- the SAS command that is used to start the server

The intermediate level of organization is called a logical server object. SAS servers of a particular type, running either on the same machine or on different machines, can be grouped into a logical server of the corresponding type. For example:

- A logical workspace server is a group of one or more workspace servers.
- A logical pooled workspace server is a group of one or more workspace servers that are configured to use server-side pooling.
- A logical stored process server is a group of one or more stored process servers.

The logical servers are then grouped into a SAS Application Server. The following figure shows a sample configuration:

Figure 4.1 SAS Application Server Components: Sample Configuration



Application servers and logical servers are logical constructs that exist only in metadata. In contrast, the server objects within a logical server correspond to actual server processes that execute SAS code.

Purpose of the Application Server Grouping

Application servers, which are groupings of logical servers, provide the following functionality in the SAS Intelligence Platform:

- SAS applications can request the services of a group of logical servers by specifying a single application server name.

For example, a user can set up SAS Data Integration Studio to send all jobs to a specific SAS Application Server. Each time the user runs a SAS Data Integration Studio job, the SAS code that is generated is executed by a workspace server that is contained in the specified application server.

- System resources, such as SAS libraries or data schemas, can be assigned to a specific application server. This assignment ensures that all of the logical servers contained in the application server will be able to access these resources as needed, subject to security restrictions.

For example, if a SAS library is assigned to a specific application server, then any application that runs jobs on that server will automatically have access to the library, subject to security restrictions.

Purpose of the Logical Server Grouping

Logical servers, which are groupings of individual servers of a specific type, provide the following functionality in the SAS Intelligence Platform:

- Users or groups of users can be granted rights to a logical server, thereby providing access to each physical server that the logical server contains. You can also use permissions to direct specific users or groups of users to specific physical servers within a logical server grouping.
- If one physical server in a logical server grouping fails, the other servers are still available to provide continuous processing.
- **Load balancing** can be implemented among the servers in a logical workspace server, logical stored process server, or logical OLAP server. Load balancing is configured by default for the SAS Pooled Workspace Server and SAS Stored Process Servers.
- **Workspace pooling** can be implemented among the servers in a logical workspace server.

For more information about SAS Application Servers, see the *SAS Intelligence Platform: Application Server Administration Guide*.

Load Balancing for SAS Workspace Servers, SAS Stored Process Servers, and SAS OLAP Servers

Load balancing is a feature that distributes work among the server processes in a logical workspace server, logical stored process server, or logical OLAP server. The load balancer runs in the object spawner, which is a program that runs on server machines, listens for incoming client requests, and launches server instances as needed. When a logical server is set up for load balancing, and the object spawner receives a client request for a server in the logical server group, the spawner directs the request to the server in the group that has the least load.

SAS Pooled Workspace Servers and SAS Stored Process Servers are load-balanced by default. For more information about load balancing, see the *SAS Intelligence Platform: Application Server Administration Guide*.

Workspace Pooling for SAS Workspace Servers

Workspace pooling creates a set of server connections that are reused. This reuse avoids the wait times that an application incurs when it creates a new server connection for each user. You can also use pooling to distribute server connections across machines.

Pooling is most useful for applications that require frequent, but brief, connections to a SAS Workspace Server. Two types of pooling are supported:

server-side pooling

a configuration in which the SAS object spawner maintains a collection of re-usable workspace server processes that are available for clients. The usage of servers in the pool is governed by authorization rules that are set on the servers in the SAS metadata. Load balancing is automatically configured for these servers.

By default, applications such as SAS Information Map Studio, the SAS Information Delivery Portal, and SAS Web Report Studio use SAS Pooled Workspace Servers to query relational data. SAS Pooled Workspace Servers are configured to use server-side pooling and load balancing.

client-side pooling

a configuration in which the client application maintains a collection of reusable workspace server processes. In releases prior to 9.2, client-side pooling was the only method of configuring pooling for workspace servers.

For more information about pooling, see the *SAS Intelligence Platform: Application Server Administration Guide*.

Logging and Monitoring for SAS Servers

The SAS Intelligence Platform uses a standard logging facility to perform logging for SAS servers. The logging facility supports problem diagnosis and resolution, performance and capacity management, and auditing and regulatory compliance. The logging facility provides the following capabilities for servers:

- Log events can be directed to multiple destinations, including files, operating system facilities, and client applications.
- For each log destination, you can configure the message layout, including the contents, the format, the order of information, and literal text.
- For each log destination, you can configure a filter to include or exclude events based on levels and message contents.
- For the metadata server, security-related events are captured, including authentication events, client connections, changes to user and group information, and permission changes.
- You can generate performance-related log events in a format that can be processed by an Application Response Measurement (ARM) 4.0 server.

To monitor servers and spawners, you can use the server management features of SAS Management Console. These features enable you to do the following:

- view server and spawner logs
- change logging levels dynamically without stopping the server
- quiesce, stop, pause, and resume servers and spawners
- validate servers and spawners and test server connections
- view information about current connections and processes
- view performance counters that provide statistics about activity that has occurred since a server or spawner was last started

Server monitoring can also be performed using third-party products for enterprise systems management.

For more information about server logging and monitoring, see the *SAS Intelligence Platform: System Administration Guide* and *SAS Logging: Configuration and Programming Reference*.

SAS Grid Computing

You can use SAS grid computing tools to manage a distributed grid environment for your SAS deployment. SAS Grid Manager, working together with Platform Suite for SAS, enables you to distribute server workloads across multiple computers on a network to obtain the following benefits:

- the ability to accelerate SAS analytical results by adding additional computers to the grid and by dividing jobs into separate processes that run in parallel across multiple servers
- the flexibility to upgrade and maintain the computing resources on which your SAS servers are deployed without disrupting operations, and to add computing resources quickly to handle increased workloads and peak demands
- continuity of service through the high availability of critical components running in the grid

Implementation of a grid environment involves planning and design efforts to determine the topology and configuration that best meets the needs of your organization. In some cases, third-party data sharing facilities or hardware load balancers might be required. For more information, see *Grid Computing in SAS* and “Introduction to Grid Computing” at <http://support.sas.com/rnd/scalability/grid/index.html>.

Chapter 5

Middle-Tier Components of the SAS Intelligence Platform

| | |
|-------------------------------------------------|-----------|
| Overview of Middle-Tier Components | 31 |
| Third-Party Software Components | 32 |
| Web Application Server | 32 |
| Java Development Kit | 33 |
| SAS Web Infrastructure Platform | 33 |
| SAS Content Server | 34 |
| SAS Foundation Services | 35 |

Overview of Middle-Tier Components

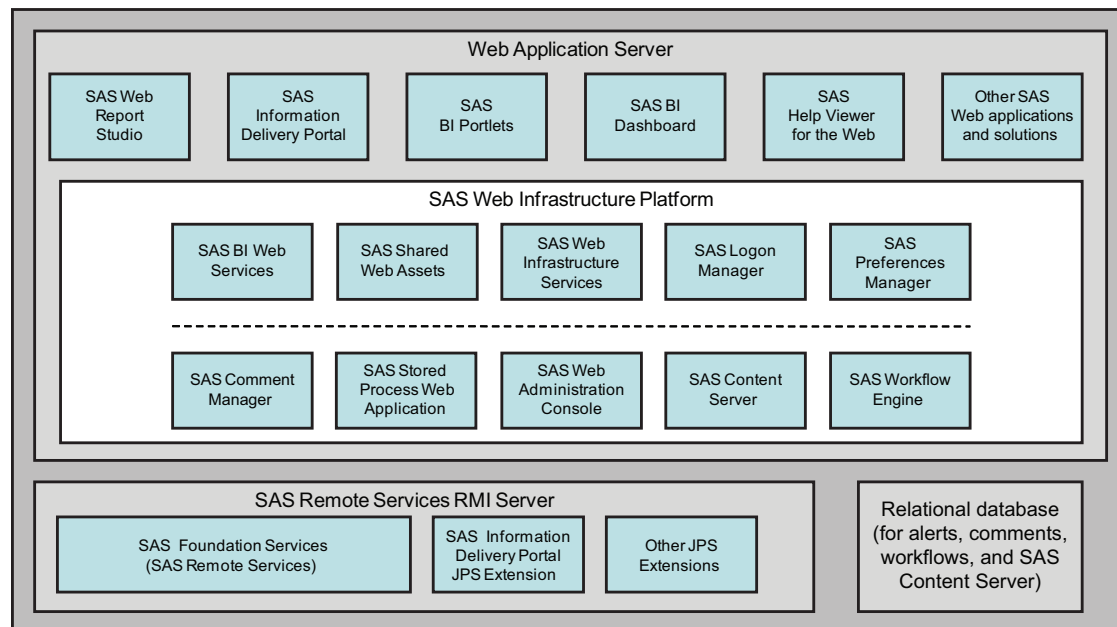
The middle tier of the SAS Intelligence Platform provides an environment in which the business intelligence Web applications, such as SAS Web Report Studio and the SAS Information Delivery Portal, can execute. These products run in a Web application server and communicate with the user by sending data to and receiving data from the user's Web browser. The middle tier applications rely on servers on the SAS server tier to perform SAS processing, including data query and analysis.

The middle-tier environment includes the following third-party software and SAS software elements:

- a Web application server
- the Java Development Kit (JDK)
- SAS Web applications, which can include SAS Web Report Studio, the SAS Information Delivery Portal, SAS BI Portlets, the SAS BI Dashboard, the SAS Help Viewer for the Web, and other SAS products and solutions
- the SAS Web Infrastructure Platform, which includes the SAS Content Server and other infrastructure applications and services
- a Java remote method invocation (RMI) server, which enables access to SAS Foundation Services and associated extension services

The following figure illustrates the middle-tier components:

Figure 5.1 Middle Tier Components



Third-Party Software Components

Web Application Server

The Web application server provides the execution environment for Web applications that run on the middle tier. The following third-party servers are supported:

- JBoss Application Server
- IBM WebSphere Application Server
- IBM WebSphere Application Server Network Deployment
- Oracle WebLogic Server

For information about the currently supported versions of these products, see “Third Party Software for SAS 9.3 Foundation” at support.sas.com/thirdpartysupport.

The following applications and services run in the Web application server environment:

- applications and services that are part of the SAS Web Infrastructure Platform.
- the SAS Web Report Studio, SAS Information Delivery Portal, SAS BI Portlets, SAS BI Dashboard, and SAS Help Viewer for the Web applications. For a description of these applications, see “Clients in the SAS Intelligence Platform” on page 37.

Depending on which products and solutions you have purchased, your site might have additional Web applications.

Java Development Kit

The Java Development Kit (JDK) is required to run the Web application server. For information about the currently supported versions of the JDK, see “Third Party Software for SAS 9.3 Foundation” at support.sas.com/thirdpartysupport.

SAS Web Infrastructure Platform

The SAS Web Infrastructure Platform is a collection of services and applications that provide common infrastructure and integration features to be used by SAS Web applications. These services and applications provide the following benefits:

- consistency in installation, configuration, and administration tasks for Web applications
- greater consistency in users' interactions with Web applications
- integration among Web applications as a result of the ability to share common resources

The following services and applications are included in the SAS Web Infrastructure Platform:

SAS BI Web Services

can be used to enable your custom applications to invoke and obtain metadata about SAS Stored Processes. Web services enable distributed applications that are written in different programming languages and that run on different operating systems to communicate using standard Web-based protocols. The most common protocol is the Simple Object Access Protocol (SOAP).

SAS Shared Web Assets

contains graph applet JARs that are shared across SAS Web applications. They display graphs in stored processes and in the SAS Stored Process Web application.

SAS Web Infrastructure Services

provide a common infrastructure for SAS Web applications. The infrastructure supports activities such as auditing, authentication, configuration, status and monitoring, e-mail, theme management, alert registration and notification, collaboration, and data sharing across SAS Web applications. The infrastructure also includes Web service interfaces that enable Java and .NET desktop applications to access the authentication service. SAS Enterprise Miner and other analytic products and solutions use these services for user authentication.

SAS Logon Manager

provides a common user authentication mechanism for SAS Web applications. It displays a dialog box for user ID and password entry, authenticates the user, and launches the requested application. SAS Logon Manager supports a single sign-on authentication model. When this model is enabled, it provides access to a variety of computing resources (including servers and Web pages) during the application session without repeatedly prompting the user for credentials.

You can configure SAS Logon Manager to display custom messages and to specify whether a logon dialog box is displayed when users log off.

In addition, you can use third-party products in conjunction with SAS Logon Manager to enable users to access multiple Web applications within the same browser session.

SAS Preferences Manager

provides a common mechanism for managing preferences for SAS Web applications. The feature enables administrators to set default preferences for locale, theme, alert notification, and time, date, and currency display. Within each Web application, users can view the default settings and update their individual preferences.

SAS Comment Manager

enables users to create comments related to business intelligence objects. Users can then reply to, search for, retire, delete, or add attachments to existing comments. This feature enables the capture of human intelligence and supports collaborative decision making related to business data.

SAS Stored Process Web Application

executes stored processes on behalf of a Web client and return results to a Web browser. The SAS Stored Process Web application is similar to the SAS/IntrNet Application Broker, and has the same general syntax and debug options.

SAS Web Administration Console

provides features for monitoring and administering middle-tier components. This browser-based interface enables administrators to do the following:

- view a list of users who are logged on to SAS Web applications, send e-mail to active users, and manage users' logins during system maintenance by preventing, terminating, and restarting user access
- view audit reports of logon and logoff activity
- create, delete, and manage permissions for folders on the SAS Content Server
- view configuration information for each middle-tier component
- dynamically adjust logging levels for some Web applications

SAS Content Server

stores digital content (such as documents, reports, and images) that is created and used by SAS Web applications. For more information, see [“SAS Content Server” on page 34](#).

SAS Workflow Engine

enables the deployment and execution of workflows and processes that are integrated with certain SAS applications and solutions. The SAS Workflow Studio client application enables business users to quickly create or modify process maps to be executed by the engine.

By default, the SAS Web Infrastructure Platform uses the SAS Framework Data Server to store alerts, comments, workflows, and SAS Content Server data. During installation, you can choose to store this data in a third-party database instead.

For more information about the SAS Web Infrastructure Platform, see the *SAS Intelligence Platform: Middle-Tier Administration Guide*.

SAS Content Server

The SAS Content Server is part of the SAS Web Infrastructure Platform. This server stores digital content (such as documents, reports, and images) that is created and used by SAS Web applications. For example, the SAS Content Server stores report definitions that are created by users of SAS Web Report Studio, as well as images and other elements that are used in reports. A process called content mapping ensures that report

content is stored using the same folder names, folder hierarchy, and permissions that the SAS Metadata Server uses to store corresponding report metadata.

In addition, the SAS Content Server stores documents and other files that are to be displayed in the SAS Information Delivery Portal or in SAS solutions.

To interact with the SAS Content Server, client applications use WebDAV-based protocols for access, versioning, collaboration, security, and searching. Administrative users can use the browser-based SAS Web Management Console to create, delete, and manage permissions for folders on the SAS Content Server.

For more information about the SAS Content Server, see the *SAS Intelligence Platform: Middle-Tier Administration Guide*.

SAS Foundation Services

SAS Foundation Services is a set of core middleware infrastructure services that integrate distributed applications on the middle tier with other components of the SAS Intelligence Platform. This suite of Java-based APIs provides the following services:

- client connections to SAS application servers
- dynamic service discovery
- user authentication
- profile management
- session management
- activity logging
- metadata and content repository access
- connection management
- WebDAV service

Extension services for information publishing, event management, and SAS Stored Process execution are also provided. All of the SAS Web applications that are described in this document use SAS Foundation Services.

For more information about SAS Foundation Services, see the *SAS Foundation Services: Administrator's Guide*.

Chapter 6

Clients in the SAS Intelligence Platform

| | |
|-----------------------------------------------------|-----------|
| Overview of Clients | 37 |
| SAS Add-In for Microsoft Office | 38 |
| SAS BI Dashboard | 39 |
| SAS BI Portlets | 39 |
| SAS Data Integration Studio | 40 |
| SAS Enterprise Guide | 40 |
| SAS Enterprise Miner | 41 |
| SAS Forecast Studio | 41 |
| SAS Information Delivery Portal | 42 |
| SAS Information Map Studio | 42 |
| JMP | 42 |
| SAS Management Console | 43 |
| SAS Model Manager | 44 |
| SAS OLAP Cube Studio | 44 |
| SAS Web Parts for Microsoft SharePoint | 45 |
| SAS Web Report Studio | 45 |
| SAS Workflow Studio | 46 |
| SAS Help Viewer for the Web | 46 |

Overview of Clients

SAS Intelligence Platform clients include desktop applications and Web-based applications. The following table lists the clients by type:

Table 6.1 SAS Intelligence Platform Clients

| Desktop Applications | Web Applications |
|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| Java Applications: <ul style="list-style-type: none"> • SAS Data Integration Studio • SAS Enterprise Miner • SAS Forecast Studio • SAS Information Map Studio • SAS Management Console • SAS Model Manager • SAS OLAP Cube Studio • SAS Workflow Studio Windows applications:* <ul style="list-style-type: none"> • SAS Add-In for Microsoft Office • SAS Enterprise Guide • JMP | <ul style="list-style-type: none"> • SAS BI Portlets • SAS BI Dashboard • SAS Help Viewer for the Web • SAS Information Delivery Portal • SAS Web Parts for Microsoft SharePoint • SAS Web Report Studio |

* JMP is also available on Macintosh and Linux.

The desktop applications are supported only on Microsoft Windows systems. The exception is SAS Management Console, which runs on all platforms except z/OS. All of the Java desktop applications require the Java Runtime Environment (JRE), which includes a Java Virtual Machine (JVM) that executes the application and a set of standard Java class libraries. If you have installed the SAS Foundation on a host, the JRE will already be present on that machine. Otherwise, you can install the JRE that is supplied by SAS before you install the first Java client.

The Web-based applications reside and execute on the middle tier (see “[Middle-Tier Components of the SAS Intelligence Platform](#)” on page 31). They require only a Web browser to be installed on each client machine, with the addition of Adobe Flash Player for SAS BI Dashboard. The applications run on a Web application server and communicate with the user by sending data to and receiving data from the user's browser. For example, an application of this type displays its user interface by sending an HTML document to the user's browser. The user can submit input to the application by sending it an HTTP response—usually by clicking a link or submitting an HTML form.

SAS Add-In for Microsoft Office

The SAS Add-In for Microsoft Office enables you to harness the power of SAS analytics from Microsoft Word, Microsoft Excel, Microsoft PowerPoint, and Microsoft Outlook. You can use the SAS Add-In for Microsoft Office for the following tasks:

- dynamically run stored processes and SAS tasks.
- analyze SAS or Excel data by using SAS tasks.
- embed the results of your stored processes and SAS tasks in your documents, spreadsheets, and presentations.

- access and view SAS data sources or any data source that is available from your SAS server. There is no size limit on the SAS data sources that you can open.
- send results from one Microsoft Office application to another. For example, you can send an analysis result that was created in Microsoft Excel to Microsoft Word or PowerPoint; and you can send reports from Microsoft Outlook to Excel, Word, or PowerPoint. When sending results, you can specify whether the results can be refreshed or whether they should be static.
- refresh SAS results to reflect changes that were made to a data source that is saved on a server or changes that were made to the SAS task or stored process.
- share results by sending e-mail, scheduling meetings, or assigning tasks in Microsoft Outlook.

For more information, see the Help for the SAS Add-In for Microsoft Office. For information about administrative tasks associated with SAS Add-In for Microsoft Office, see the *SAS Intelligence Platform: Desktop Application Administration Guide*.

SAS BI Dashboard

The SAS BI Dashboard enables users to create, maintain, and view dashboards to monitor key performance indicators that convey how well an organization is performing. The application is Web-based, leveraging the capabilities of Adobe Flash Player, and can be accessed from within the SAS Information Delivery Portal.

The SAS BI Dashboard includes an easy-to-use interface for creating dashboards that include graphics, text, colors, and hyperlinks. Dashboards can link to SAS reports and analytical results, SAS Strategic Performance Management scorecards and objects, externally generated data, and virtually anything that is addressable by a Uniform Resource Identifier (URI).

All content is displayed in a role-based, secure, customizable, and extensible environment. End users can customize how information appears on their personal dashboards.

For more information, see the SAS BI Dashboard Help, which is available from within the product, and the *SAS BI Dashboard: User's Guide*, available at <http://support.sas.com>. For information about administrative tasks associated with the SAS BI Dashboard, see the *SAS Intelligence Platform: Web Application Administration Guide*.

SAS BI Portlets

SAS BI Portlets enable users to access, view, and work with business intelligence content that is stored on the SAS Metadata Server and the SAS Content Server. The portlets are seamlessly integrated into the SAS Information Delivery Portal and provide added value to portal users. The portlets include the SAS Collection Portlet, the SAS Diagnostics Portlet, the SAS Navigator Portlet, the SAS Report Portlet, and the SAS Stored Process Portlet. For information about configuring, deploying, and administering SAS BI Portlets, see the *SAS Intelligence Platform: Web Application Administration Guide*.

SAS Data Integration Studio

SAS Data Integration Studio is a visual design tool that enables you to consolidate and manage enterprise data from a variety of source systems, applications, and technologies. This software enables you to create process flows that accomplish the following tasks:

- extract, transform, and load data for use in data warehouses and data marts
- cleanse, migrate, synchronize, replicate, and promote data for applications and business services

SAS Data Integration Studio enables you to create metadata that defines sources, targets, and the processes that connect them. This metadata is stored in one or more shareable repositories. SAS Data Integration Studio uses the metadata to generate or retrieve SAS code that reads sources and creates targets in physical storage. Other applications that share the same repositories can use the metadata to access the targets and use them as the basis for reports, queries, or analyses.

Through its metadata, SAS Data Integration Studio provides a single point of control for managing the following resources:

- data sources (from any platform that is accessible to SAS and from any format that is accessible to SAS)
- data targets (to any platform that is accessible to SAS, and to any format that is supported by SAS)
- processes that specify how data is extracted, transformed, and loaded from a source to a target
- jobs that organize a set of sources, targets, and processes (transformations)
- source code generated by SAS Data Integration Studio
- user-written source code, including legacy SAS programs

Note: SAS Data Integration Studio was formerly named SAS ETL Studio.

For more information, see the SAS Data Integration Studio Help, which is available from within the product, and the *SAS Data Integration Studio: User's Guide*, which is available at <http://support.sas.com>. For information about administrative tasks associated with SAS Data Integration Studio, see the *SAS Intelligence Platform: Desktop Application Administration Guide*.

SAS Enterprise Guide

SAS Enterprise Guide is a project-oriented application that is designed to enable quick access to much of the analytic power of SAS software for statisticians, business analysts, and SAS programmers. SAS Enterprise Guide provides the following functionality:

- a point-and-click user interface to SAS servers
- transparent data access to both SAS data and other types of data
- interactive task windows that lead you through dozens of analytical and reporting tasks
- a program editor with autocomplete and integrated syntax help

- the ability to use the highest quality SAS graphics
- the ability to export results to other Windows applications and the Web
- the ability to schedule your project to run at a later time

SAS Enterprise Guide also enables you to create SAS Stored Processes and to store that code in a repository that is available to a SAS Stored Process Server. (Stored processes are SAS programs that are stored on a server and are executed by client applications.) Stored processes are used for Web reporting and analytics, among other things.

For more information, see the SAS Enterprise Guide Help, which is available from within the product. For information about administrative tasks associated with SAS Enterprise Guide, see the *SAS Intelligence Platform: Desktop Application Administration Guide*.

SAS Enterprise Miner

SAS Enterprise Miner is an example of the SAS analytics products that use the SAS Intelligence Platform. This application streamlines the data mining process to create highly accurate predictive and descriptive models based on the analysis of vast amounts of data from across an enterprise. SAS Enterprise Miner enables you to develop data mining flows that consist of data exploration, data transformations, and modeling tools to predict or classify business outcomes.

For more information, see the SAS Enterprise Miner Help, which is available from within the product. For information about administrative tasks associated with SAS Enterprise Miner, see the *SAS Intelligence Platform: Desktop Application Administration Guide*.

SAS Forecast Studio

SAS Forecast Studio is the client component of the SAS Forecast Server, which is an example of the SAS analytics products that use the SAS Intelligence Platform. The SAS Forecast Server is a large-scale, automatic forecasting solution that enables organizations to produce huge quantities of high-quality forecasts quickly and automatically. SAS Forecast Studio provides a graphical interface (based on Java) to the forecasting and time series analysis procedures contained in SAS High-Performance Forecasting and SAS/ETS software.

SAS High-Performance Forecasting automatically selects the appropriate model for each item being forecast, based on user-defined criteria. Holdout samples can be specified so that models are selected not only by how well they fit past data but by how appropriate they are for predicting the future. If the best forecasting model for each item is unknown or the models are outdated, a maximum level of automation can be chosen in which all three forecasting steps (model selection, parameter estimation, and forecast generation) are performed. If suitable models have been determined, you can keep the current models and reestimate the model parameters and generate forecasts. For maximum processing speed, you can keep previously selected models and parameters and choose to simply generate the forecasts.

For more information, see the *SAS High-Performance Forecasting: User's Guide* and the *SAS/ETS User's Guide*.

SAS Information Delivery Portal

The SAS Information Delivery Portal enables you to aggregate data from a variety of sources and present the data in a Web browser. The Web browser content might include the output of SAS Stored Processes, links to Web addresses, documents, syndicated content from information providers, SAS Information Maps, SAS reports, and Web applications.

Using the portal, you can distribute different types of content and applications as appropriate to internal users, external customers, vendors, and partners. You can use the portal along with the Publishing Framework to publish content to SAS publication channels or WebDAV repositories, to subscribe to publication channels, and to view packages published to channels. The portal's personalization features enable users to organize information on their desktops in a way that makes sense to them.

For more information, see the SAS Information Delivery Portal Help, which is available from within the product. For information about administrative and development tasks associated with the SAS Information Delivery Portal, see the *SAS Intelligence Platform: Web Application Administration Guide* and *Developing Portlets for the SAS Information Delivery Portal*, which are available at <http://support.sas.com>.

SAS Information Map Studio

SAS Information Map Studio enables data modelers and data architects to create and manage SAS Information Maps, which are business metadata about your data. Information maps enable you to surface your data warehouse data in business terms that typical business users understand, while storing key information that is needed to build appropriate queries.

Information maps provide the following benefits:

- Users are shielded from the complexities of the data.
- Data storage is transparent to the consumers of information maps, regardless of whether the underlying data is relational or multidimensional, or whether the data is in a SAS data set or in a third-party database system.
- Business formulas and calculations are predefined, which makes them usable on a consistent basis.
- Users can easily query data for answers to business questions without having to know query languages.

For more information, see the SAS Information Map Studio Help, which is available from within the product, and *SAS Information Map Studio: Getting Started with SAS Information Maps*, which is available at <http://support.sas.com>.

JMP

JMP is interactive, exploratory data analysis and modeling software for the desktop. JMP makes data analysis—and the resulting discoveries—visual and helps to

communicate those discoveries to others. JMP presents results both graphically and numerically. By linking graphs to each other and to the data, JMP makes it easier to see the trends, outliers, and other patterns that are hidden in your data.

Through interactive graphs that link statistics and data, JMP offers analyses from the basic (univariate descriptive statistics, ANOVA, and regression) to the advanced (generalized linear, mixed, and nonlinear models, data mining, and time series models). A unified, visual approach makes these techniques available to all levels of users.

The JMP Scripting Language can be used to create interactive applications and to communicate with SAS. JMP reads text files, Microsoft Excel spreadsheets, SAS data sets, and data from any ODBC data source.

JMP is included with SAS Visual BI, which can be ordered with SAS Enterprise BI Server.

SAS Management Console

SAS Management Console provides a single point of control for administering your SAS servers and for managing metadata objects that are used throughout the SAS Intelligence Platform. You can use SAS Management Console to connect to the SAS Metadata Server and view and manage the metadata objects that are stored in the server's metadata repositories.

SAS Management Console uses an extensible plug-in architecture, enabling you to customize the console to support a wide range of administrative capabilities. The plug-ins that are provided with the application enable you to manage the following resources:

- server definitions
- library definitions
- user, group, and definitions
- resource access controls
- metadata repositories
- job schedules

The plug-ins also enable you to monitor server activity, configure and run metadata server backups, view metadata interdependencies, and stop, pause, and resume servers. Only certain users can view and use plug-ins. A user's access to plug-ins depends on the which roles the user is assigned to and which capabilities are assigned to those roles.

In SAS Management Console, you can also view the folders that store metadata for business intelligence content such as libraries, tables, jobs, cubes, information maps, and reports. The folders include a personal folder for each individual user. For shared data, you can create a customized folder structure and assign appropriate permissions at each level in the folder hierarchy. Import and export wizards enable you to easily move or promote content from one part of the folder tree to another, or from one metadata server to another.

Note: Other client applications also display folders if they contain content that is relevant to the application and that the user has permission to view.

For more information, see the SAS Management Console Help, which is available from within the product. For detailed information about specific tasks that you can perform in SAS Management Console, see the administration guides for the SAS Intelligence Platform at <http://support.sas.com/93administration>.

SAS Model Manager

SAS Model Manager is an example of the SAS analytics products that use the SAS Intelligence Platform. This product provides a central model repository and model management environment for predictive and analytical models. SAS Model Manager also provides procedural templates that facilitate model and metadata management across organizational areas, throughout the life cycle of the model.

In SAS Model Manager, a model is associated with properties that describe it, such as the type of analysis, algorithm, input and output variables, quality measures, reports, attribute rankings, and target scores. These properties are stored in the model metadata. Model metadata can also document conformance to industry requirements, such as best practice information and credit scoring audit trails.

SAS Model Manager provides the following functionality:

- stores models in a central secure repository
- manages and promotes models through the analytic process using the SAS Workflow or SAS Model Manager life cycle
- compares and validates models
- performs scoring tests for model assessment
- publishes the model life cycle and performance data over established reporting channels
- provides model performance monitoring and dashboard reports for models deployed to production
- provides the ability to retrain models based on recent data

SAS OLAP Cube Studio

SAS OLAP Cube Studio enables you to design and create online analytical processing (OLAP) cubes, register cube metadata in a SAS Metadata Repository, and save physical cube data in a specified location. Using the application's Cube Designer Wizard, you can specify the following:

- the data source for a cube
- the cube design and architecture
- measures of the cube for future queries
- initial aggregations for the cube

SAS OLAP Cube Studio also includes functions and wizards that enable you to create additional cube aggregations, to add calculated members to a cube, and to modify existing calculated members.

For more information, see the SAS OLAP Cube Studio Help, which is available from within the product. For information about administrative tasks associated with SAS OLAP Cube Studio, see the *SAS Intelligence Platform: Desktop Application Administration Guide*.

SAS Web Parts for Microsoft SharePoint

SAS Web Parts for Microsoft SharePoint enables you to add SAS content directly to your Web site through the use of Microsoft Windows SharePoint Services. This application makes the following SAS Web Parts available from the Web Parts Gallery:

SAS BI Dashboard Web Part

The SAS BI Dashboard Web Part enables you to include interactive dashboards that have been created in SAS BI Dashboard on a SharePoint page. If an indicator in the dashboard is linked to other SAS content, users can click the indicator to view the additional SAS content. For more information, see “[SAS BI Dashboard](#)” on page 39.

SAS Stored Process Web Part

The SAS Stored Process Web Part enables you to include a SAS Stored Process on a SharePoint page. Users can then run the stored process and view the results on the page.

For more information, see the SAS Web Parts for Microsoft Sharepoint Help. For information about administrative tasks associated with SAS Web Parts for Microsoft Sharepoint, see the *SAS Intelligence Platform: Web Application Administration Guide*.

SAS Web Report Studio

SAS Web Report Studio enables you to create, view, and organize reports. You can use SAS Web Report Studio for the following tasks:

Creating reports

Beginning with a simple and intuitive view of your data provided by SAS Information Maps (created in SAS Information Map Studio), you can create reports based on either relational or multidimensional data sources. Advanced users can access tables and cubes directly, without the need for information maps.

You can use the Report Wizard to quickly create simple reports, or you can use Edit mode to create sophisticated reports that have multiple data sources, each of which can be filtered. These reports can include various combinations of list tables, crosstabulation tables, graphs, images, and text. Using Edit mode, you can adjust the style to change colors and fonts. You can also insert stored processes, created by business analysts who are proficient in SAS, that contain instructions for calculating analytical results. The stored process results can be rendered as part of a report or as a complete report.

Additional layout and query capabilities are available for advanced users. These users can incorporate custom calculations and complex filter combinations, multiple queries, prompts, and SAS analytical results into a single document. In addition, advanced users can use headers, footers, images, and text to incorporate corporate standards, confidentiality messages, and hyperlinks into reports. Advanced users can also access tables and cubes directly.

Viewing and working with reports

While viewing reports, you can filter, sort, and rank the data that is shown in list tables, crosstabulation tables, and graphs. With multidimensional data, you can drill down on data in crosstabulation tables and graphs, and drill through to the underlying data.

Organizing reports

You can create folders and subfolders for organizing your reports. Information consumers can use keywords to find the reports that they need. Reports can be shared with others or kept private. You can schedule reports to run at specified times and distribute them via e-mail, either as PDF attachments or as embedded HTML files. You can also publish reports to one or more publication channels.

Printing and exporting reports

You can preview a report in PDF and print the report, or save and e-mail it later. You have control over many printing options, including page orientation and page range. You can also export data as a spreadsheet and export graphs as images. You can also export data for list tables, crosstabulation tables, and graphs. The output can be viewed in Microsoft Excel or Microsoft Word.

For more information, see the SAS Web Report Studio Help and the *SAS Web Report Studio: User's Guide*, which are available from within the product. For information about administrative tasks associated with SAS Web Report Studio, see the *SAS Intelligence Platform: Web Application Administration Guide*.

SAS Workflow Studio

SAS Workflow Studio enables you to modify and extend the default workflows and process definitions that are provided with certain SAS applications and solutions. SAS Workflow Studio is a business process-modeling tool for the rapid development of process maps that reflect an application's business logic. It enables business users to build, organize, and deploy complex business processes, and to adjust processes efficiently when business demands change. The process layer can be integrated with any number of existing applications.

For more information, see the SAS Workflow Studio Help.

SAS Help Viewer for the Web

The SAS Help Viewer for the Web enables users to view and navigate SAS Online Help in Web applications such as the SAS Information Delivery Portal, the SAS BI Dashboard, and SAS Web Report Studio. For information about administration tasks associated with SAS Help Viewer for the Web, see the *SAS Intelligence Platform: Web Application Administration Guide*.

Chapter 7

Security Overview

| | |
|--------------------------------------------------------------|-----------|
| Overview of SAS Intelligence Platform Security | 47 |
| Authorization and Permissions Overview | 48 |
| Metadata-Based Authorization | 48 |
| Multiple Authorization Layers | 49 |
| Roles and Capabilities Overview | 49 |
| Authentication and Identity Management Overview | 49 |
| Single Sign-On in the SAS Intelligence Platform | 50 |
| Single Sign-On for SAS Desktop Applications | 50 |
| Single Sign-On for SAS Web Applications | 50 |
| Single Sign-On for Data Servers and Processing Servers | 51 |
| Overview of Initial Users | 51 |
| Encryption Overview | 51 |
| Security Reporting and Logging Overview | 52 |

Overview of SAS Intelligence Platform Security

The security features of the SAS Intelligence Platform offer the following benefits:

- single sign-on from and across disparate systems
- secure access to data and metadata
- role-based access to application features
- confidential transmission and storage of data
- logging and auditing of security events
- access control reporting

The SAS Intelligence Platform's security model cooperates with external systems such as the host environment, the Web realm, and third-party databases. To coordinate identity information, SAS keeps a copy of one or more IDs (such as host, Active Directory, LDAP, or Web account IDs) for each user. This requirement does not apply to any users for whom a generic PUBLIC identity is sufficient.

For a comprehensive discussion of security and detailed information about security administration activities, see the *SAS Intelligence Platform: Security Administration Guide*, which is available at <http://support.sas.com/93administration>.

Authorization and Permissions Overview

Metadata-Based Authorization

Authorization is the process of determining which users have which permissions for which resources. The SAS Intelligence Platform includes an authorization mechanism that consists of access controls that you define and store in a metadata repository. These metadata-based controls supplement protections from the host environment and other systems. You can use the metadata authorization layer to manage access to the following resources:

- almost any metadata object (for example, reports, data definitions, information maps, jobs, stored processes, and server definitions)
- OLAP data
- relational data (depending on the method by which the data is accessed)

You can set permissions at several levels of granularity:

- Repository-level controls provide default access controls for objects that have no other access controls defined.
- Resource-level controls manage access to a specific item such as a report, an information map, a stored process, a table, a column, a cube, or a folder. The controls can be defined individually (as explicit settings) or in patterns (by using access control templates).
- Fine-grained controls affect access to subsets of data within a resource. You can use these controls to specify who can access particular rows within a table or members within a cube dimension.

You can assign permissions to individual users or to user groups. Each SAS user has an identity hierarchy that starts with the user's individual SAS identity and can include multiple levels of nested group memberships.

The effect of a particular permission setting is influenced by any related settings that have higher precedence. For example, if a report inherits a grant from its parent folder but also has an explicit denial, the explicit setting determines the outcome.

The available metadata-based permissions are summarized in the following table.

Table 7.1 *Metadata-Based Permissions*

| Permissions | Use |
|-------------------------------------------------------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| ReadMetadata, WriteMetadata, WriteMemberMetadata, CheckInMetadata | Use to control user interactions with a metadata object. |
| Read, Write, Create, or Delete | Use to control user interactions with the underlying computing resource that is represented by a metadata object; and to control interactions with some metadata objects, such as dashboard objects. |

| Permissions | Use |
|-------------|-----------------------------------------------------------------------------------------------------------------------------------------|
| Administer | Use to control administrative interactions (such as starting or stopping) with the SAS server that is represented by a metadata object. |

Multiple Authorization Layers

A user's ability to perform a particular action is determined not only by metadata-based access controls but also by external authorization mechanisms such as operating system permissions and database controls. To perform a particular action, the user must have the necessary permissions in all of the applicable authorization layers. For example, regardless of the access controls that have been defined for the user in the metadata repository, the user cannot access a particular file if the operating system permissions do not permit the action.

Roles and Capabilities Overview

The SAS implementation of roles enables administrators to manage the availability of application features such as menu items, plug-ins, and buttons. Applications that have roles include the SAS Add-In for Microsoft Office, SAS Enterprise Guide, SAS Forecast Studio, SAS Management Console, and SAS Web Report Studio. For example, role memberships determine whether a user can see the Server Manager plug-in (in SAS Management Console), compare data (in SAS Enterprise Guide), or directly open an information map (in SAS Web Report Studio). Administrators can assign roles to users and to groups.

An application feature that is under role management is called a capability. Each application that supports roles provides a fixed set of explicit and implicit capabilities. Explicit capabilities can be incrementally added to or removed from any role (other than the unrestricted role, which always provides all explicit capabilities). An implicit capability is permanently bound to a certain role. A contributed capability is an implicit or explicit capability that is assigned through role aggregation. If one role is designated as a contributing role for another role, all of the first role's capabilities become contributed capabilities for the second role.

In general, roles are separate from permissions and do not affect access to metadata or data.

Authentication and Identity Management Overview

Authentication is an identity verification process that attempts to determine whether users (and other entities) are who they say they are. In the simplest case, users already have accounts that are known to the metadata server's host. For example, if the metadata server is on UNIX, then users might have accounts in an LDAP provider that the UNIX host recognizes. If the metadata server is on Windows, then users might have Active Directory accounts.

For accountability, we recommend creating an individual SAS identity for each person who uses the SAS environment. These identities enable administrators to make access distinctions and audit individual actions in the metadata layer. The identities also provide personal folders for each user. The metadata server maintains its own copy of each user ID for the purpose of establishing a SAS identity.

Identity management tasks can be performed manually using SAS Management Console or by using the following batch processes:

- To load user information into the metadata repository, you first extract user and group information from one or more enterprise identity sources. Then you use SAS bulk-load macros to create identity metadata from the extracted information. SAS provides sample applications that extract user and group information and logins from an Active Directory server and from UNIX `/etc/passwd` and `/etc/group` files.
- To periodically update user information in the metadata repository, you extract user and group information from your enterprise identity sources and from the SAS metadata. Then you use SAS macros to compare the two sets of data and identify the needed updates. After validating the changes, you use SAS macros to load the updates into the metadata repository.

Note: You cannot use these batch processes to manage passwords. Users can manage their own passwords with the SAS Personal Login Manager.

The metadata identity information is used by the security model's credential management and authorization features. For example, when a user logs on to SAS Data Integration Studio, the metadata server wants to know who the user is so that it can determine which libraries, stored processes, and jobs should be displayed in the desktop client. If a user makes a request in SAS Data Integration Studio to run a job against an Oracle table, the Oracle server wants to know who the user is so that it can determine whether the user has access to the data in the table.

Single Sign-On in the SAS Intelligence Platform

Single Sign-On for SAS Desktop Applications

For desktop applications such as SAS Information Map Studio, SAS Enterprise Guide, SAS Data Integration Studio, SAS OLAP Cube Studio, and SAS Management Console, you can use the following single sign-on features:

- You can enable Integrated Windows authentication so that users will not receive a logon prompt when they launch applications. Integrated Windows authentication is a Microsoft technology that generates and validates Windows identity tokens. All participating clients and servers must authenticate against the same Windows domain (or against domains that trust one another).
- Users can also avoid the initial logon prompt by selecting the option to save their credentials in a connection profile. (This option can be disabled on a site-wide basis.)

Single Sign-On for SAS Web Applications

You can enable Web authentication so that users will not receive a logon prompt when they launch SAS Web applications such as SAS Web Report Studio and the SAS Information Delivery Portal. In this configuration, SAS Web applications use whatever authentication scheme you have set up in your Web environment. For example, if your

Web environment is integrated with a third-party authentication provider, then the SAS Web applications participate in that scheme.

Single Sign-On for Data Servers and Processing Servers

Seamless access to SAS Stored Process Servers, SAS OLAP Servers, SAS Content Servers, and SAS Pooled Workspace Servers is provided through SAS token authentication. This mechanism causes participating SAS servers to accept users who are connected to the metadata server. No individual external accounts are required, no user passwords are stored in the metadata, and no reusable credentials are transmitted.

Seamless access to SAS Workspace Servers can be provided through SAS token authentication, Integrated Windows authentication, or credential reuse. With credential reuse, when a user provides credentials in the initial logon dialog box, the credentials are added to the user's in-memory credential cache (user context) and then retrieved when access to the workspace server is required.

You can also use Integrated Windows authentication to provide direct connections to OLAP servers (for example, from a data provider) when there is no active connection to the metadata server.

Overview of Initial Users

During installation, several initial user accounts are created. Some of these accounts are created for all installations, some are optional, and some are created only if certain software components are installed. The required users include the following:

- The SAS Administrator account and the SAS Trusted User Account. These users are generally set up as internal accounts, which exist in metadata but are not known to the host machine. The SAS Administrator account has access to all metadata, regardless of SAS permissions settings. The SAS Trusted User is a privileged service account that can act on behalf of other users when connecting to the metadata server.
- The SAS Spawned Servers account and the SAS Installer account, which must be defined in the operating system of certain server machines. The SAS Spawned Servers account is the initially configured process owner for pooled workspace servers and stored process servers. The SAS Installer Account is used to install and configure SAS software. On UNIX and z/OS systems, this account is also the owner of configuration directories and their contents and is the process owner for items such as the metadata server, the OLAP server, and the object spawner.

Other initial users include the LSF Administrator and LSF User, which are required if Platform Suite for SAS is installed. In addition, the SAS Anonymous Web Service User is an optional account that is used to grant clients access to applicable SAS Web Infrastructure Platform components. Most installations set up this user as an internal account, which exists in metadata but is not known to the host machine.

Encryption Overview

SAS offers encryption features to help you protect information on disk and in transit. When passwords must be stored, they are encrypted or otherwise encoded. Passwords that are transmitted by SAS are also encrypted or encoded. You can choose to encrypt all traffic instead of encrypting only credentials.

If you have SAS/SECURE, you can use an industry standard encryption algorithm such as AES. SAS/SECURE is an add-on product that is licensed separately and offers maximum protection. This product includes support of the Federal Information Processing Standard (FIPS) 140-2 encryption specification. If you do not have SAS/SECURE, you can use the SASProprietary algorithm to help protect information.

Security Reporting and Logging Overview

Security reporting creates a snapshot of metadata layer access control settings. SAS provides the %MDSECDS autocall macro to enable you to easily build data sets of permissions information. You can use those data sets as the data source for security reports. You can also identify changes in settings by comparing data sets that are generated at different times.

Security logging records security-related events as part of the system-wide logging facility. The events that are captured include authentication events, client connections, changes to user and group information, and permission changes. For more information about logging, see [“Logging and Monitoring for SAS Servers”](#) on page 28 and the *SAS Logging: Configuration and Programming Reference*.

Glossary

access control template

a reusable named authorization pattern that you can apply to multiple resources. An access control template consists of a list of users and groups and indicates, for each user or group, whether permissions are granted or denied. Short form: ACT.

ACT

See access control template.

Application Response Measurement

the name of an application programming interface that was developed by an industry partnership and which is used to monitor the availability and performance of software applications. ARM monitors the application tasks that are important to a particular business. Short form: ARM.

ARM

See Application Response Measurement.

authentication

See client authentication.

authentication provider

a software component that is used for identifying and authenticating users. For example, an LDAP server or the host operating system can provide authentication.

authorization

the process of determining which users have which permissions for which resources. The outcome of the authorization process is an authorization decision that either permits or denies a specific action on a specific resource, based on the requesting user's identity and group memberships.

browser

See Web browser.

capability

an application feature that is under role-based management. Typically, a capability corresponds to a menu item or button. For example, a Report Creation capability might correspond to a New Report menu item in a reporting application. Capabilities are assigned to roles.

client authentication

the process of verifying the identity of a person or process for security purposes.

client-side pooling

a configuration in which the client application maintains a collection of reusable workspace server processes.

credentials

the user ID and password for an account that exists in some authentication provider.

cube

See OLAP cube.

cube

See multidimensional database.

data mart

a subset of the data in a data warehouse. A data mart is optimized for a specific set of users who need a particular set of queries and reports.

data set

See SAS data set.

data warehouse

a collection of data that is extracted from one or more sources for the purpose of query, reporting, and analysis. Data warehouses are generally used for storing large amounts of data that originates in other corporate applications or that is extracted from external data sources.

database management system

a software application that enables you to create and manipulate data that is stored in the form of databases. Short form: DBMS.

database server

a server that provides relational database services to a client. Oracle, DB/2 and Teradata are examples of relational databases.

DBMS

See database management system.

encryption

the act or process of converting data to a form that is unintelligible except to the intended recipients.

foundation services

See SAS Foundation Services.

HTTP

a protocol for transferring data to the Internet. HTTP provides a way for servers and Web clients to communicate. It is based on the TCP/IP protocol.

HyperText Transfer Protocol

See HTTP.

identity

See metadata identity.

information map

a collection of data items and filters that provides a user-friendly view of a data source. When you use an information map to query data for business needs, you do not have to understand the structure of the underlying data source or know how to program in a query language.

Integrated Object Model

the set of distributed object interfaces that make SAS software features available to client applications when SAS is executed as an object server. Short form: IOM.

Integrated Object Model server

See IOM server.

Integrated Windows authentication

a Microsoft technology that facilitates use of authentication protocols such as Kerberos. In the SAS implementation, all participating components must be in the same Windows domain or in domains that trust each other.

internal account

a SAS account that you can create as part of a user definition. Internal accounts are intended for metadata administrators and some service identities; these accounts are not intended for regular users.

IOM

See Integrated Object Model.

IOM server

a SAS object server that is launched in order to fulfill client requests for IOM services.

IWA

See Integrated Windows authentication.

Java Development Kit

See JDK.

Java RMI

See remote method invocation.

Java Virtual Machine

See JVM.

JDK

a software development environment that is available from Oracle Corporation. The JDK includes a Java Runtime Environment (JRE), a compiler, a debugger, and other tools for developing Java applets and applications.

JVM

a program that interprets Java programming code so that the code can be executed by the operating system on a computer. The JVM can run on either the client or the server. The JVM is the main software component that makes Java programs portable across platforms. A JVM is included with JDKs and JREs from Oracle Corporation, as well as with most Web browsers.

LDAP

a protocol that is used for accessing directories or folders. LDAP is based on the X.500 standard, but it is simpler and, unlike X.500, it supports TCP/IP.

Lightweight Directory Access Protocol

See LDAP.

load balancing

for IOM bridge connections, a program that runs in the object spawner and that uses an algorithm to distribute work across object server processes on the same or separate machines in a cluster.

logical server

the second-level object in the metadata for SAS servers. A logical server specifies one or more of a particular type of server component, such as one or more SAS Workspace Servers.

MDDB

See multidimensional database.

metadata

descriptive data about data that is stored and managed in a database, in order to facilitate access to captured and archived data for further use.

metadata identity

a metadata object that represents an individual user or a group of users in a SAS metadata environment. Each individual and group that accesses secured resources on a SAS Metadata Server should have a unique metadata identity within that server.

metadata LIBNAME engine

the SAS engine that processes and augments data that is identified by metadata. The metadata engine retrieves information about a target SAS data library from metadata objects in a specified metadata repository.

metadata object

a set of attributes that describe a table, a server, a user, or another resource on a network. The specific attributes that a metadata object includes vary depending on which metadata model is being used.

metadata server

a server that provides metadata management services to one or more client applications. A SAS Metadata Server is an example.

multidimensional database

a specialized data storage structure in which data is presummarized and cross-tabulated and then stored as individual cells in a matrix format, rather than in the row-and-column format of relational database tables. The source data can come either from a data warehouse or from other data sources. MDDBs can give users quick, unlimited views of multiple relationships in large quantities of summarized data.

object spawner

a program that instantiates object servers that are using an IOM bridge connection. The object spawner listens for incoming client requests for IOM services. When the spawner receives a request from a new client, it launches an instance of an IOM server to fulfill the request. Depending on which incoming TCP/IP port the request

was made on, the spawner either invokes the administrator interface or processes a request for a UUID (Universal Unique Identifier).

OLAP

See online analytical processing.

OLAP cube

a logical set of data that is organized and structured in a hierarchical, multidimensional arrangement to enable quick analysis of data. A cube includes measures, and it can have numerous dimensions and levels of data.

OLAP schema

a container for OLAP cubes. A cube is assigned to an OLAP schema when it is created, and an OLAP schema is assigned to a SAS OLAP Server when the server is defined in the metadata. A SAS OLAP Server can access only the cubes that are in its assigned OLAP schema.

online analytical processing

a software technology that enables users to dynamically analyze data that is stored in multidimensional database tables (cubes).

online analytical processing

a software technology that enables users to dynamically analyze data that is stored in multidimensional database tables (cubes).

parallel I/O

a method of input and output that takes advantage of multiple CPUs and multiple controllers, with multiple disks per controller to read or write data in independent threads.

parallel processing

a method of processing that divides a large job into several smaller jobs that can be executed in parallel on multiple CPUs.

permission

the type of access that a user or group has to a resource. The permission defines what the user or group can do with the resource. Examples of permissions are ReadMetadata and WriteMetadata.

plug-in

a file that modifies, enhances, or extends the capabilities of an application program. The application program must be designed to accept plug-ins, and the plug-ins must meet design criteria specified by the developers of the application program.

RDBMS

See relational database management system.

relational database management system

a database management system that organizes and accesses data according to relationships between data items. The main characteristic of a relational database management system is the two-dimensional table. Examples of relational database management systems are DB2, Oracle, Sybase, and Microsoft SQL Server.

remote method invocation

a Java programming feature that provides for remote communication between programs by enabling an object that is running in one Java Virtual Machine (JVM) to

invoke methods on an object that is running in another JVM, possibly on a different host. Short form: RMI.

RMI

See remote method invocation.

role

See user role.

SAS Application Server

a logical entity that represents the SAS server tier, which in turn comprises servers that execute code for particular tasks and metadata objects.

SAS Content Server

a server that stores digital content (such as documents, reports, and images) that is created and used by SAS client applications. To interact with the server, clients use WebDAV-based protocols for access, versioning, collaboration, security, and searching.

SAS data set

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

SAS Foundation Services

a set of core infrastructure services that programmers can use in developing distributed applications that are integrated with the SAS platform. These services provide basic underlying functions that are common to many applications. These functions include making client connections to SAS application servers, dynamic service discovery, user authentication, profile management, session context management, metadata and content repository access, activity logging, event management, information publishing, and stored process execution.

SAS Framework Data Server

a database server that is the default location for middle-tier data such as alerts, comments, and workflows, as well as data for the SAS Content Server and SAS Service Parts Optimization. The server is provided as an alternative to using a third-party DBMS. The server cannot be used as a general-purpose data store.

SAS Management Console

a Java application that provides a single user interface for performing SAS administrative tasks.

SAS Metadata Repository

a container for metadata that is managed by the SAS Metadata Server.

SAS Metadata Server

a multi-user server that enables users to read metadata from or write metadata to one or more SAS Metadata Repositories.

SAS OLAP Server

a SAS server that provides access to multidimensional data. The data is queried using the multidimensional expressions (MDX) language.

SAS Open Metadata Architecture

a general-purpose metadata management facility that provides metadata services to SAS applications. The SAS Open Metadata Architecture enables applications to exchange metadata, which makes it easier for these applications to work together.

SAS SPD Server

a SAS Scalable Performance Data Server. An SPD Server restructures data in order to enable multiple threads, running in parallel, to read and write massive amounts of data efficiently.

SAS Stored Process

a SAS program that is stored on a server and defined in metadata, and which can be executed by client applications. Short form: stored process.

SAS Stored Process Server

a SAS IOM server that is launched in order to fulfill client requests for SAS Stored Processes.

SAS token authentication

a process in which the metadata server generates and verifies SAS identity tokens to provide single sign-on to other SAS servers. Each token is a single-use, proprietary software representation of an identity.

SAS Web Infrastructure Platform

a collection of middle-tier services and applications that provide infrastructure and integration features that are shared by SAS Web applications and other HTTP clients.

SAS Workspace Server

a SAS IOM server that is launched in order to fulfill client requests for IOM workspaces.

SAS/CONNECT server

a server that provides SAS/CONNECT services to a client. When SAS Data Integration Studio generates code for a job, it uses SAS/CONNECT software to submit code to remote computers. SAS Data Integration Studio can also use SAS/CONNECT software for interactive access to remote libraries.

SAS/SHARE server

the result of an execution of the SERVER procedure, which is part of SAS/SHARE software. A server runs in a separate SAS session that services users' SAS sessions by controlling and executing input and output requests to one or more SAS libraries.

SASProprietary algorithm

a fixed encoding algorithm that is included with Base SAS software. The SASProprietary algorithm requires no additional SAS product licenses. It provides a medium level of security.

server-side pooling

a configuration in which a SAS object spawner maintains a collection of reusable workspace server processes that are available for clients. The usage of servers in this pool is governed by the authorization rules that are set on the servers in the SAS metadata.

service

one or more application components that an authorized user or application can call at any time to provide results that conform to a published specification. For example, network services transmit data or provide conversion of data in a network, database services provide for the storage and retrieval of data in a database, and Web services interact with each other on the World Wide Web.

single sign-on

an authentication model that enables users to access a variety of computing resources without being repeatedly prompted for their user IDs and passwords. For example, single sign-on can enable a user to access SAS servers that run on different platforms without interactively providing the user's ID and password for each platform. Single sign-on can also enable someone who is using one application to launch other applications based on the authentication that was performed when the user initially logged on.

SMP

See symmetric multiprocessing.

spawner

See object spawner.

SPD Server

See SAS SPD Server.

SSO

See single sign-on.

stored process

See SAS Stored Process.

symmetric multiprocessing

a hardware and software architecture that can improve the speed of I/O and processing. An SMP machine has multiple CPUs and a thread-enabled operating system. An SMP machine is usually configured with multiple controllers and with multiple disk drives per controller. Short form: SMP.

theme

a collection of specifications (for example, colors, fonts, and font styles) and graphics that control the appearance of an application.

thread

a single path of execution of a process that runs on a core on a CPU.

transformation

in data integration, an operation that extracts data, transforms data, or loads data into data stores.

user role

a set of permissions that define which actions a user, or a group of users, can take in an application.

warehouse

See data warehouse.

Web application

an application that is accessed via a Web browser over a network such as the Internet or an intranet. SAS Web applications are Java Enterprise Edition (JEE) applications that are delivered via Web application archive (WAR) files. The applications can depend on Java and non-Java Web technologies.

Web authentication

a configuration in which users of Web applications and Web services are verified at the Web perimeter and the metadata server trusts that verification.

Web browser

a software application that is used to view Web content, and also to download or upload information. The browser submits URL (Uniform Resource Locator) requests to a Web server and then translates the HTML code into a visual display.

Web-distributed authoring and versioning

a set of extensions to the HTTP protocol that enables users to collaboratively edit and manage files on remote Web servers. Short form: WebDAV.

WebDAV

See Web-distributed authoring and versioning.

Index

A

- access controls 24, 48
- accessibility features 2
- aggregating data 42
- analytical models 7, 44
- analytics 3, 6
- application servers 26
 - functionality of 26
 - purpose of grouping 26
- architecture 9
- Aster nCluster databases 16
- authentication 49
 - Integrated Windows 51
 - token 51
- authorization 48, 50
 - metadata-based 48
 - multiple layers 49

B

- batch servers 22
- business intelligence 3, 5

C

- CA-IDMS 16
- capabilities 49
- clickstream data 11
- client-side pooling 28
- clients 10, 14
 - by type 37
 - JMP 42
 - overview of 37
 - SAS Add-In for Microsoft Office 38
 - SAS BI Dashboard 39
 - SAS BI Portlets 39
 - SAS Data Integration Studio 40
 - SAS Enterprise Guide 40
 - SAS Enterprise Miner 41
 - SAS Forecast Studio 41
 - SAS Help Viewer for the Web 46

- SAS Information Delivery Portal 42
- SAS Information Map Studio 42
- SAS Management Console 43
- SAS Model Manager 44
- SAS OLAP Cube Studio 44
- SAS Web Parts for Microsoft SharePoint 45
- SAS Web Report Studio 45
- SAS Workflow Studio 46
- content mapping 34
- credential reuse 51
- cube metadata 44
- cubes 5, 18, 22, 44

D

- dashboards 5, 39
- data management 2, 3
 - DataFlux Data Management Platform 3
 - SAS Data Integration Studio 3
 - SAS Data Quality Server 4
 - SAS Data Surveyor 4
 - SAS Metadata Repository 7
 - SAS/ACCESS 4
- data mining 6, 41
- data servers 22
 - single sign-on for 51
- data sets 16
 - as default SAS storage 16
 - data storage 4
- data source design applications 19
- data sources 9, 10
 - metadata management for 18
- DATA step batch servers 22
- data storage 4, 15
 - cubes 18
 - data sets 16
 - default SAS storage 16
 - high-performance computing 4, 16
 - in-database 4, 16
 - multidimensional databases 18

- multidimensional storage 4
 - parallel storage 5, 17
 - relational storage 4
 - third-party databases 4, 16
 - databases
 - Aster nCluster 16
 - DB2 16
 - Greenplum 16
 - hierarchical 16
 - multidimensional 18
 - Netezza 17
 - network model systems 16
 - Oracle 17
 - relational 16
 - Teradata 17
 - third-party 4, 11, 16
 - DataFlux Data Management Platform 3
 - DB2 databases 16
 - default SAS storage 16
 - desktop applications 37
 - single sign-on for 50
 - digital content 34
 - documentation
 - SAS Help Viewer for the Web 46
- E**
- encryption 51
 - exporting reports 46
- F**
- Federal Processing Standard (FIPS) 140-2
 - encryption 52
 - fine-grained controls 48
 - FIPS 140-2 encryption 52
 - forecasting 6, 41
- G**
- generic batch servers 22
 - Greenplum databases 16
 - grid computing 29
 - Grid Management Services 7
- H**
- Help
 - SAS Help Viewer for the Web 46
 - hierarchical databases 16
 - high-performance computing 16
- I**
- IBM WebSphere Application Server 32
 - Network Deployment 32
 - identity management 49
 - automated support for 49
 - identity verification 49
 - in-database technology 4, 16
 - information maps 5, 42
 - initial users 51
 - Integrated Windows authentication 51
- J**
- Java batch servers 22
 - Java Development Kit 5 33
 - JBoss Application Server 32
 - JMP 7, 42
 - JMP Scripting Language 43
- L**
- load balancing 27
 - Load Sharing Facility (LSF) 7
 - logging
 - for SAS servers 28
 - for security 52
 - logical pooled workspace servers 25
 - logical server objects 25
 - logical servers
 - purpose of grouping 27
 - logical stored process servers 25
 - logical workspace servers 25
 - Logon Manager 33
 - LSF Administrator account 51
 - LSF User account 51
- M**
- managing data
 - See [data management](#)
 - metadata
 - creating and administering 24
 - cube metadata 44
 - in SAS Metadata Repository 23
 - managing data sources in 18
 - organization of 25
 - metadata LIBNAME engine 19
 - metadata objects 23
 - managing 43
 - metadata repository
 - See [SAS Metadata Repository](#)
 - metadata-based authorization 48
 - metadata-based permissions 48
 - middle tier 10, 13
 - components 31
 - third-party software components 32
 - models, analytical 7, 44

- monitoring SAS servers 28
- multidimensional databases 18
- multidimensional storage 4
- multiple authorization layers 49

N

- Netezza databases 17
- network model database systems 16

O

- object spawner 22
 - load balancing and 27
 - monitoring 28
 - SAS Spawned Servers account 51
- OLAP cubes 5, 18, 22, 44
- online Help
 - SAS Help Viewer for the Web 46
- Oracle databases 17
- Oracle WebLogic Server 32

P

- parallel storage 5, 17
 - how it works 17
 - options for implementing 17
- passwords 50, 51
- permissions 48
 - metadata-based 48
- Platform Suite for SAS 7
- plug-ins 43
- pooled workspace server 13, 22
 - logical 25
- pooling
 - client-side 28
 - server-side 13, 22, 27
 - workspace 27
- portlets 5, 39
- Preferences Manager 34
- printing reports 46
- process flows 40
- Process Manager 7
- processing servers
 - single sign-on for 51

R

- relational databases 16
- relational storage 4
- reporting, security 52
- reports 45
 - creating 45
 - organizing 46
 - printing and exporting 46
 - viewing and working with 45

- repository-level controls 48
- resource-level controls 48
- roles 24, 49

S

- SAS Add-In for Microsoft Office 5, 38
- SAS Administrator account 51
- SAS Anonymous Web Service User account 51
- SAS BI Dashboard 5, 39
- SAS BI Dashboard Web Part 45
- SAS BI Portlets 5, 39
- SAS BI Server 2
- SAS BI Web Services for Java 33
- SAS Business Analytics Framework 1
- SAS Comment Manager 34
- SAS Content Server 22, 34
- SAS Data Integration Server 2
- SAS Data Integration Studio 3, 40
- SAS Data Quality Server 4
- SAS Data Surveyor 4, 11
 - for clickstream data 11
 - for SAP 11
- SAS Enterprise BI Server 2
- SAS Enterprise Data Integration Server 2
- SAS Enterprise Guide 5, 40
- SAS Enterprise Miner 6, 41
- SAS Forecast Server 6, 41
- SAS Forecast Studio 41
- SAS Foundation Services 35
- SAS Framework Data Server 11, 22
- SAS grid computing 29
- SAS Grid Manager 7, 29
- SAS Help Viewer for the Web 46
- SAS identity 24, 49
- SAS In-Database 4, 16
- SAS Information Delivery Portal 5, 42
- SAS Information Map Studio 5, 42
- SAS Installer account 51
- SAS Intelligence Platform 1
 - accessibility features 2
 - architecture 9
 - components 2
 - roles in 49
 - security 47
 - servers 21
 - single sign-on 50
 - strategic benefits 7
- SAS Logon Manager 33
- SAS Management Console 7, 43
 - defining data source objects 19
- SAS Metadata Repository 7
 - loading user information into 50
 - metadata in 23
 - updating user information in 50

- SAS Metadata Server 12, 21, 23
 - controlling system access 24
 - creating and administering metadata 24
 - metadata in SAS Metadata Repository 23
 - SAS Model Manager 7, 44
 - SAS object spawner 13, 22
 - load balancing and 27
 - monitoring 28
 - SAS Spawned Servers account 51
 - SAS OLAP Cube Studio 5, 44
 - SAS OLAP Server 12, 18, 22
 - data storage 4
 - SAS Pooled Workspace Server 13, 22
 - SAS Preferences Manager 34
 - SAS servers 9, 11, 21
 - logging and monitoring for 28
 - SAS Content Server 34
 - SAS Metadata Server 12
 - SAS OLAP Server 12
 - SAS Pooled Workspace Server 13
 - SAS Stored Process Server 13
 - SAS Workspace Server 12
 - SAS Shared Web Assets 33
 - SAS Spawned Servers account 51
 - SAS SPD Engine
 - data storage 5
 - parallel storage 17
 - symmetric multiprocessing (SMP) 17
 - SAS SPD Server
 - data storage 5
 - parallel storage 17
 - SAS Stored Process Server 13, 22
 - load balancing 27
 - logical 25
 - SAS Stored Process Web Application 34
 - SAS Stored Process Web Part 45
 - SAS tables
 - See [data sets](#)
 - SAS token authentication 51
 - SAS Trusted User Account 51
 - SAS Web Administration Console 34
 - SAS Web Infrastructure Platform 33
 - SAS Web Infrastructure Platform Services 33
 - SAS Web Parts for Microsoft SharePoint 45
 - SAS Web Report Studio 5, 45
 - SAS Workflow Engine 34
 - SAS Workflow Studio 46
 - SAS Workspace Server 12, 21
 - load balancing 27
 - logical 25
 - logical pooled 25
 - pooled 13, 22
 - workspace pooling 27
 - SAS/ACCESS 4
 - SAS/CONNECT servers 22
 - security 47
 - authentication and identity management 49
 - authorization and permissions 48
 - encryption 51
 - initial users 51
 - reporting and logging 52
 - roles 49
 - single sign-on 50
 - server objects 25
 - server-side pooling 13, 22, 27
 - servers 9, 11, 21
 - single sign-on 50
 - for data servers and processing servers 51
 - for desktop applications 50
 - for Web applications 50
 - SMP (symmetric multiprocessing) 17
 - spawner 22
 - load balancing and 27
 - monitoring 28
 - SAS Spawned Servers account 51
 - storage
 - See [data storage](#)
 - See [parallel storage](#)
 - stored process servers
 - See [SAS Stored Process Server](#)
 - stored processes 13, 22, 25, 27, 41
 - symmetric multiprocessing (SMP) 17
 - system access
 - SAS Metadata Server control of 24
- T**
- tables
 - See [data sets](#)
 - Teradata databases 17
 - third-party databases 11, 16
 - data storage 4
 - third-party software
 - middle-tier components 32
 - threaded reads 4
 - token authentication 51
- U**
- updating
 - user information in metadata repository 50
 - user information
 - loading into metadata repository 50
 - updating in metadata repository 50
 - users, initial 51

V

verifying identities [49](#)

W

Web application server [32](#)

Web applications [13](#)

 single sign-on for [50](#)

Web-based applications [37](#)

workspace pooling [27](#)

workspace servers

See [SAS Workspace Server](#)

