SAS® Environment Manager 2.5 Administration

User's Guide

SAS® Documentation

The correct bibliographic citation for this manual is as follows: SAS Institute Inc. 2015. *SAS® Environment Manager 2.5 Administration: User's Guide*. Cary, NC: SAS Institute Inc.

SAS® Environment Manager 2.5 Administration: User's Guide

Copyright © 2015, SAS Institute Inc., Cary, NC, USA

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

For a hard-copy 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 License Rights; Restricted Rights: The Software and its documentation is commercial computer software developed at private expense and is provided with RESTRICTED RIGHTS to the United States Government. Use, duplication or disclosure of the Software by the United States Government is subject to the license terms of this Agreement pursuant to, as applicable, FAR 12.212, DFAR 227.7202-1(a), DFAR 227.7202-3(a) and DFAR 227.7202-4 and, to the extent required under U.S. federal law, the minimum restricted rights as set out in FAR 52.227-19 (DEC 2007). If FAR 52.227-19 is applicable, this provision serves as notice under clause (c) thereof and no other notice is required to be affixed to the Software or documentation. The Government's rights in Software and documentation shall be only those set forth in this Agreement.

SAS Institute Inc., SAS Campus Drive, Cary, North Carolina 27513-2414.

August 2015

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 trademarks of their respective companies.

Contents

Chapter 1 • Using SAS Environment Manager Administration **1**
What is SAS Environment Manager Administration? 1
Starting SAS Environment Manager Administration 1
Finding Your Way Around 2
Controlling Access to Modules 3

Chapter 2 • Managing User Access **5**
Overview of User Administration 6
Add a User 6
Add an Administrator 7
Add a Custom Group 7
Add a Custom Role 8
Assign Members to a Group or Role 8
Update the Stored Password in a Login 9
Delete an Identity 9
Store Contact Information 9
Store DBMS Credentials 10
Determine Whether a User Has an Internal Account 10
Adjust Policies for an Internal Account 11
Change a Role's Capabilities 11
About User Administration 11

Chapter 3 • Managing Metadata Access **21**
Overview of Metadata Administration 21
Apply an ACT 22
Create an ACT 22
Update an ACT 23
Add an Explicit Grant or Denial 24
About Metadata Administration 24
Icons in Access Management 30
Permission Origins 32
Best Practices for Permissions 34

Chapter 4 • Managing Servers **37**
Using the Servers Module 37
Create a New Server Definition 37
Filtering the Server List 38

Chapter 5 • Managing Libraries **39**
Using the Libraries Module 39
Creating a New Library Definition 39
Registering a Table 40
Change Assigned SAS Servers 41

Chapter 6 • Managing Backups **43**
Overview of SAS Backup Manager 43
Viewing Backup and Recovery History 47
Running an Unscheduled Backup 50
Viewing and Modifying the Backup Policy 51
Viewing and Modifying the Backup Schedule 53

Recommended Reading **55**
Index **57**

Chapter 1

Using SAS Environment Manager Administration

What is SAS Environment Manager Administration? 1
Starting SAS Environment Manager Administration 1
Finding Your Way Around 2
Controlling Access to Modules 3

What is SAS Environment Manager Administration?

SAS Environment Manager Administration is a web-based administration solution for a SAS environment that enables you to manage SAS resources and resource definitions, including:

- folders and objects
- authorization controls
- user and user group definitions
- library definitions
- database server definitions
- SAS content backups

The application provides these functions through the use of modules. Each module manages a specific type of SAS metadata. In addition to the modules that are provided to manage these types of core metadata, other modules are available to manage a specialized type of metadata or provide specialized management functions.

Starting SAS Environment Manager Administration

To access SAS Environment Manager Administration, use your web browser to go to http://<*localhost*>:7080, where *localhost* is the machine on which the SAS Environment Manager server is installed. The SAS Login Manager prompts you for your credentials

and the application starts. After you specify your credentials in the SAS Login Manager, the main SAS Environment Manager application appears. Click **Administration** in the menu bar to open SAS Environment Manager Administration.

Finding Your Way Around

When you start SAS Environment Manager Administration, the application displays the Folders module. This view enables you to view and manage SAS folders and the metadata objects that they contain.

To switch to a different module, click the **Side Menu** icon ☰ to open the **Side Menu**, which displays a list of all of the available modules. Click on a module name to open it and view the specific objects that the module manages.

As you open object definitions in the modules, the object counter icon

in the toolbar keeps track of the definitions that are open and provides easy access to an open definition. The counter on the icon indicates the number of object definitions that are open. Click on the icon to display a menu of all open definitions. Select an item in the menu to go to that definition. An asterisk beside an entry in the menu indicates that the definition has been changed but not yet saved.

When you are viewing a table of object definitions in a module, you can click on a column heading to sort or filter the table. Clicking the **Options** icon ⋮ in a table enables you to select the columns that you want to display.

Object definitions open on the Basic Properties page. The title of the page is displayed at the top of the page, next to the entry's name. To view other property pages for the definition, click on the page title to display a menu of the page titles.

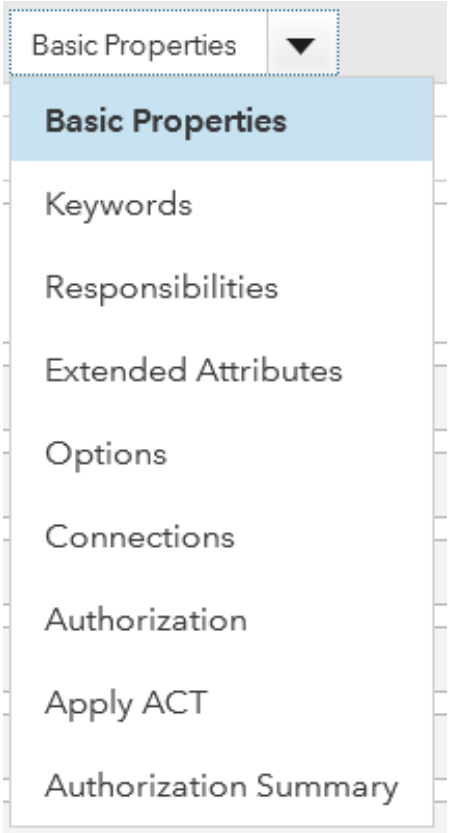

Note: If you are using Microsoft Internet Explorer 11 and you change your computer's locale setting, you must clear the browser's cache to view the SAS Environment Manager menus in the locale's language.

Controlling Access to Modules

You can use metadata role capabilities to control access to the SAS Environment Manager Administration modules. Removing the capability for a module prevents the module from appearing in the SAS Environment Manager Administration side menu. Any user to whom the role applies will not be able to access the module. For information about modifying capabilities, see "About Roles and Capabilities" on page 13.

To control access to modules, follow these steps:

1. Start SAS Management Console and select the User Manager plug-in.
2. Select only the **Show Roles** check box to display the defined roles.
3. Open the **Management Console: Content Management** role.
4. Select the **Capabilities** tab.
5. Expand the **Applications** ⇨ **Management Console 9.4** ⇨ **Plug-ins** folders.
6. Modify these capabilities as needed:

   Data Library Manager
   controls access to the Libraries module

Folder View
controls access to the Folders module

Server Manager
controls access to the Servers module

User Manager
controls access to the Users module

Chapter 2

Managing User Access

Overview of User Administration 6

Add a User 6

Add an Administrator 7

Add a Custom Group 7

Add a Custom Role 8

Assign Members to a Group or Role 8

Update the Stored Password in a Login 9

Delete an Identity 9

Store Contact Information 9

Store DBMS Credentials 10
- Store Shared Credentials 10
- Store Individual Credentials 10

Determine Whether a User Has an Internal Account 10

Adjust Policies for an Internal Account 11

Change a Role's Capabilities 11

About User Administration 11
- About User Administration 11
- About Users 12
- About SAS Administrators 12
- About Groups 13
- About Roles and Capabilities 13
- About Members 14
- About Logins 14
- About Internal Accounts 16
- About Authentication Domains 17
- About Passwords 18
- About External Identities 18
- Requirement: Unique Names and IDs 19

Overview of User Administration

This chapter addresses user administration in the metadata layer, which is provided by the SAS Metadata Server.

The SAS Environment Manager Users module supports some of the user administration tasks that are provided by the User Manager plug-in to SAS Management Console, including the following:

- creation and maintenance of users, groups, and roles
- management of group and role memberships
- management of logins and internal accounts
- assignment of capabilities to roles

In addition, this application provides the following enhanced user administration features, including the following:

- When memberships are displayed, all memberships are shown, including indirect and implicit relationships.
- For each role, capability origins information explains how that role gets each of its capabilities (directly and from contributing roles).

Note: This application provides improved access to information. It does not introduce any changes to the underlying security model.

Many of the general features of SAS Environment Manager are useful in user administration. For example, you can open multiple objects. Each object is displayed in its own page.

Note: Changes that you make to one object are not reflected in other open objects until those changes are saved.

Add a User

1. In the Users Module, click and select **New User**.
2. On the New User window, enter a name and optional display name. Then click **Save**. The new user object opens on the Basic Properties page.
3. From the drop-down menu, select **Accounts**. On the Accounts page, click to add a login.
   - In the Domain column, select **DefaultAuth**.
   - In the Stored User ID column, enter the user's external account ID. You can use any account (LDAP, Active Directory, host, or other) that is known to the metadata server's host.

     Note: For a Windows account, qualify the ID (for example, *WIN\myID* or *myID@mycompany.com*).

If your site uses web authentication, refer to the following table for adapted instructions:

Table 2.1 *Adapted Instructions for Sites That Use Web Authentication*

Type of User	Adapted Instructions*
Someone who uses only web applications	Select the web realm authentication domain (such as **web**) instead of **DefaultAuth**, and enter the user's web realm ID.
Someone who uses both web and desktop applications	Complete the standard instructions and also add a web realm login.

* If the web user IDs and the metadata server user IDs are identical, and the web applications do not use a standard workspace server, you do not need to follow these adapted instructions.

4. (Optional) From the drop-down menu, select **Member of**. On the Member of page, click to make a user a direct member of one or more groups and roles.

   Note: The user automatically belongs to PUBLIC (everyone who can access the metadata server) and SASUSERS (those members of PUBLIC who have a well-formed user definition).

5. (Optional) To provide seamless access to a third-party server such as a DBMS, either give the user a second login or make the user a member of a group that has a shared login for the third-party server.
6. Save the new user.

Note: Do not use the settings in any of the Authorization: pages to modify authorizations for a user. The settings in this window have no effect on what the user can do. They affect the permissions for the user object itself.

Note: If the user accesses a standard workspace server with Windows host credentials, make sure the user has the "Log on as a batch job" Windows privilege on that host. The user's Windows account should be a member of a Windows group that is named SAS Server Users or something similar.

Add an Administrator

To add an administrator, follow the instructions for adding a user. On the **Member of** page, click . In the Direct Memberships window, move the SAS Administrators group to the **Direct member of** list box. The new user becomes a member of the SAS Administrators group.

Add a Custom Group

1. In the Users Module, click and select **New Group**.
2. In the New Group window, enter a name. A display name and repository are optional. Then, click **Save**. The new group object opens to the Basic Properties page.

Note: In general, groups are in the foundation repository. If you have custom repositories, you can create groups in those repositories too.

3. From the drop-down menu, select **Members**. On the Members page, click to directly assign members to the new group.
4. If you want to make the new group a member of other groups or roles, use the **Member of** page.

   Note: The ability to nest group dependencies is limited. Suppose you make a group a member of another group. As soon as a branch in the tree returns to an item that is previously listed at any point in the branch, SAS Environment Manager stops the nested dependencies in order to avoid an infinite loop. For example, suppose you have three groups named Group1, Group2, and Group3. If Group1 is a member of Group2 and Group2 is a member of Group3, then you cannot make Group3 a member of both Group1 and Group2.
5. If you are using this group to make a shared account available, add a shared login on the Accounts page. See "Logins for Groups" on page 15.
6. Save the new group.

Note: Do not use the settings in any of the Authorization: pages to modify authorizations for a user. The settings on these pages have no effect on what the user can do.

Add a Custom Role

1. In the Users Module, click and select **New Role**.
2. In the New Role window, enter a name. A display name and repository are optional. Then, click **Save**. The new role object opens in a tab.

   Note: In general, roles are in the foundation repository. If you have custom repositories, you can create roles in those repositories too.
3. From the drop-down menu, select **Members**. On the Members page, click to directly assign members to the new role.
4. Save the new role.

Note: Do not use the settings in any of the Authorization: pages to modify authorizations for a user. The settings in this window have no effect on what the user can do.

Assign Members to a Group or Role

1. In the Users module, open the group or role that you want to update.
2. On the **Members** page, click to directly assign members.
3. Save the group or role.

Update the Stored Password in a Login

1. Open the user or group whose external password has changed.
2. From the drop-down menu, select **Accounts**. On the Accounts page, click the cell that you need to update.

   Note: Logins are visible only if you have user administration capabilities, you are looking at your own user definition, or you are looking at the definition for a group that you belong to.
3. In the Store Password (Optional) window, enter and confirm the new password.
4. Save the user or group.

Delete an Identity

CAUTION:
When you delete a user, group, or role, you lose all of that identity's metadata associations. Creating a new identity with the same name does not restore the associations.

1. In the Users module, locate the identity that you want to delete.
2. Right-click the identity and select **Delete**.
3. In the confirmation message box, click **Yes**.

Store Contact Information

Note: You cannot store contact information for groups or roles.

1. Open the user that you want to update.
2. On the **Contact Information** page, make changes as needed.
3. Save the user.

TIP If you batch synchronize users and want to preserve contact information that you enter interactively, use a consistent value in the **Type** field. In your synchronization code, you can use this value to define exceptions that exclude this data from the batch update.

Store DBMS Credentials

Store Shared Credentials

Note: These instructions apply to third-party servers that do not accept the credentials with which users initially log on to SAS clients. These instructions are also appropriate for providing seamless access to other servers that require credentials that are different from the credentials with which users initially log on to SAS clients.

1. In the Users module, open (or create) the group that you will use to manage the shared DBMS account. For example, if you want all users to share the account, use the PUBLIC group.
2. On the Accounts page, click ➕ to add a login.
   - In the Domain column, select the authentication domain for the DBMS.

     Note: In the DBMS server metadata, the authentication domain is specified on the connection object.
   - In the Stored User ID column, enter the DBMS user ID.
   - In the Stored Password (Optional) column, store the DBMS password.
3. From the drop-down menu, select Members. On the Members page, make sure that everyone who needs to use the shared account is a member.
4. Save the group.

Store Individual Credentials

Follow the instructions in the preceding topic, but add the login to a user definition instead of a group definition.

Note: If a user has more than one available login in a particular authentication domain, the login that is closest to the user is used. If there is tie, you might not be able to predict which will be used. One example of a tie is when a user is a direct member of two groups and both groups have logins in the same authentication domain.

Determine Whether a User Has an Internal Account

To determine whether a user has an internal account, examine the **Internal account** section of the user's **Accounts** page. If a user has an internal account, their internal user ID is listed in that location. In general, regular users do not have internal accounts.

Note: Internal accounts are visible only if you have user administration capabilities or you are looking at your own user definition.

Note: An administrator can enable an internal account for a user by setting the **Internal Account** option on the user’s Accounts page to **Create internal account**.

Adjust Policies for an Internal Account

1. In the Users module, open the user whose internal account policies you want to change.
2. From the drop-down menu, select **Accounts**. Enable the **Internal Account** option.
3. Make changes to one or more of the properties on the page.
4. Save your changes.

Note: The properties are displayed only when an administrator accesses the **Accounts** page for a user who has an internal account.

Change a Role's Capabilities

The initial capabilities-to-roles mapping is appropriate in many cases. If necessary, you can change the set of capabilities that a role provides.

CAUTION:
There is no automated method for reverting a role back to its original set of capabilities. Instead of adjusting the capabilities of a predefined role, consider creating a new role. Before you alter role capabilities, make sure you have a complete and current backup. For more information, see *SAS Intelligence Platform: System Administration Guide*.

1. Open the role that you want to update.
2. Use either or both of the following techniques:
   - Incrementally add or remove capabilities from the role by selecting or clearing check boxes on the Basic Capabilities page.
   - Give the role the capabilities of one or more other roles by using the Contributing Roles page.
3. Save the role.

About User Administration

About User Administration

SAS Environment Manager provides some of the capabilities that are available in SAS Management Console. SAS Environment Manager is not currently a replacement for SAS Management Console, and no functionality has been removed from SAS Management Console.

In order to make access distinctions and to track user activity, each requesting user must be identified. The main purpose of user administration is to provide information that facilitates the identification of users. In general, SAS stores one external account ID for

each user. SAS uses its copy of these IDs to establish a unique identity for each connecting user. All of a user's metadata-layer group memberships, role memberships, and permission assignments are ultimately tied to that user's SAS identity.

Note: For identification purposes, only the account IDs are needed. SAS does not store external passwords for identification purposes.

As an alternative to creating a login for a user or administrator, you can give the new user or administrator an internal account. Enable the **Internal Account** option on the user's **Accounts** page. An account is enabled when the **Internal Account** option settings are displayed. Before you add an internal account, see "Limitations of Internal Accounts" on page 16.

TIP As an alternative to interactively creating and maintaining identity information, you can write a program that performs these tasks as batch processes. See the user import macros documentation in *SAS Intelligence Platform: Security Administration Guide*.

TIP You can use the metadata promotion tools to import and export identity information. See *SAS Intelligence Platform: System Administration Guide*.

About Users

A user is an individual person or service identity.

You should create an individual SAS identity for each person who uses the SAS environment. You can then make access distinctions in the metadata layer and establish a personal folder for each user.

Note: If generic access is sufficient for some of your users, those users can instead share the generic PUBLIC group identity. Not all applications accept PUBLIC-only users.

An individual SAS identity is established by coordination between two sets of identity information:

- in an external system, a user account
- in the metadata, a user definition that includes a copy of the external account ID

To give someone an individual SAS identity, you create a metadata user definition that includes a copy of their external account ID. For example, in the simplest configuration, each user has an account that is known to the metadata server's host.

- If the metadata server is running in Windows, users typically have Active Directory accounts.
- If the metadata server is running in UNIX, users might have UNIX accounts. Sometimes a UNIX host is configured to recognize LDAP, Active Directory, or other types of accounts.

Note: For metadata administrators and some service identities, you can use a SAS internal account instead of an external account.

About SAS Administrators

When creating SAS administrators, consider the following:

- If you want to make someone an unrestricted administrator, assign them to the Metadata Server: Unrestricted role.

- Administrators should not also serve as regular users. If you want someone to be an administrator only some of the time, create two user definitions for that person.
  - One definition is based on an external account and is not a member of SAS Administrators.
  - The other definition is based on an internal account and is a member of SAS Administrators.

  A dual user logs on with an internal account in order to use administrative privileges and with an external account the rest of the time.

About Groups

A group is a set of users.

Creating groups enables you to simplify security management in the following ways:

- It is more efficient to assign permissions to groups than to individual users.
- If you need to store passwords in the metadata, you can reduce the amount of required maintenance by using a group to make one shared account available to multiple users.
- It is sometimes more efficient to manage role membership by assigning groups to roles instead of assigning users directly to roles.
- If you need to manage permissions for distinct classes of access, you can create custom groups. For example, you might create a group for each business unit or functional area of responsibility.

TIP A group's membership can include other groups as well as individual users, allowing you to create a nested group structure.

About Roles and Capabilities

A role manages the availability of application features such as menu items.

An application feature that is under role-based management is called a capability. Anyone who is a member of a role has all of that role's capabilities. Roles and capabilities have the following characteristics:

- Roles determine which user interface elements (such as menu items or modules) you see when you use an application. In general, roles do not protect data or metadata.
- Having a certain capability is not an alternative to meeting permission requirements. Permission requirements and capability requirements are cumulative.
- Roles and groups serve distinct purposes. You cannot assign permissions to a role or capabilities to a group.
- Capabilities are additive. Assigning someone to a role never reduces what that person can do.
- If necessary, you can adjust the distribution of capabilities by changing role memberships or by customizing the mapping of roles-to-capabilities.
- If you need to decrease the level of granularity, you can create a new role that aggregates two or more existing roles. For example, you might create a role that includes all capabilities other than those of the most privileged roles.

- If you need to increase the level of granularity, you can create a new role that provides only a subset of the capabilities of a predefined role.
- If you need to create a cross-application role for a particular type of functionality, you can create custom roles. For example, you might create an OLAP role that includes the OLAP capabilities from SAS Enterprise Guide and the SAS Add-In for Microsoft Office.

For information about the predefined administrative roles, see *SAS Intelligence Platform: System Administration Guide*.

About Members

A member is a user or group that is assigned to a group or role.

When adding members to a group or role, consider the following:

- You cannot use the SAS Environment Manager user interface to make a role a member of a group or of another role. You can instead make one role contribute all of its capabilities to another role.
- On a group definition, do not confuse the **Members** tab with the **Member of** tab. Use a group's **Member of** tab only if you want to make that group a member of other groups or roles.
- To reduce complexity, do not make the implicit groups (SASUSERS and PUBLIC) members of other groups. These groups can be members of certain roles.

About Logins

What is a Login?

A login is a SAS copy of information about an external account. Every login must include a user ID. In a login for a Windows account, the ID must be qualified (for example, *userID@company.com*, *domain\userID*, or *machine\userID*).

Logins are displayed as follows:

personal logins
: A user's personal logins are displayed on the **Accounts** page in the user's definition.

group logins
: A group login is assigned to a group and can be used by any member of that group. A user's group logins are not displayed on the user's **Accounts** page. You must check the properties of each group that the user belongs to in order to determine whether any of those groups have logins.

TIP The requirement to provide a qualified ID for a Windows account applies to the SAS copy of the ID. It is usually not necessary to qualify the user ID that you provide when you log on to a SAS application.

TIP If you do provide a qualified ID when you log on, you must use the same format that was used in your login. For example, Windows environments might accept both *WIN\me* and *Me.MyLastName@mycompany.com*, but SAS can understand only one of these qualified forms (the form in which the SAS copy of the ID is stored).

Logins for Users

Each user should have a login that establishes their SAS identity. You do not need to include a password in this login. The password column displays eight asterisks if a

password is stored. For example, this is how Joe's login might look when a user administrator views Joe's **Accounts** settings:

Domain	Stored User ID	Stored Password (Optional)
DefaultAuth	WIN\Joe	

A user might have additional logins that provide access to other systems. For example, if Joe has his own Oracle account, he might have these two logins:

Domain	Stored User ID	Stored Password (Optional)
DefaultAuth	WIN\Joe	
OracleAuth	ORAjoe	********

Note: The Oracle login should include a copy of Joe's Oracle password.

If a site uses web authentication, the requirements are different. For example, if Joe uses both web and desktop applications at such a site, Joe might have these three logins:

Domain	Stored User ID	Stored Password (Optional)
DefaultAuth	WIN\Joe	
OracleAuth	ORAjoe	********
web	WEB\joe	

Note: Like his DefaultAuth login, Joe's web login is used only to launch clients, so there is no need to create a SAS copy of Joe's web realm password.

Logins for Groups

Logins are not required for groups. The main reason to assign a login to a group is to make a shared account available to multiple users. A group login contains a SAS copy of the user ID and password for a shared account. For example, to provide shared access to DB2, a group might have a login that looks like this:

Domain	Stored User ID	Stored Password (Optional)
DB2Auth	sharedDB2id	********

All members of the group can see and use this login. Because this login is for a third-party database, a copy of the DBMS account password should be stored in this login.

About Internal Accounts

What is an Internal Account?

An internal account is a SAS account that the metadata server authenticates independently, without relying on an external authentication provider such as the operating system. Use internal accounts only for administrators and some service identities. For these purposes, an internal account is an acceptable substitute for an external account with a corresponding login. For example, the SAS Administrator and the SAS Trusted User can be based on internal accounts.

Benefits of Internal Accounts

Internal accounts have these advantages:

- Internal accounts provide an alternative to creating external accounts for SAS internal purposes such as inter-process communication.
- Internal accounts can be maintenance free. You do not have to synchronize internal accounts with another user registry.
- Internal accounts are usable only in the SAS realm, so they reduce exposure to the rest of your security environment.

Limitations of Internal Accounts

Although you can add an internal account to any user definition, internal accounts are not intended for regular users. Someone who has only an internal account cannot perform the following tasks:

- seamlessly launch a standard workspace server that runs under their own individual identity
- participate in Integrated Windows authentication or web authentication
- add, delete, initialize, or unregister a foundation repository

Policies for Internal Accounts

By initial policy, these server-level settings are in effect:

- Accounts do not expire and are not suspended due to inactivity.
- Passwords must be at least six characters in length, do not have to include mixed case or numbers, and do not expire.
- The five most recent passwords for an account cannot be reused for that account.
- There is no mandatory time delay between password changes.
- For an account that has a password expiration period, there is a forced password change on first use and after the password is reset by someone other than the account owner. By initial policy, passwords do not expire, so there are no forced password changes.

These default policies are set in the metadata server's omaconfig.xml file. For more information, see *SAS Intelligence Platform: Security Administration Guide*.

About Authentication Domains

What is an Authentication Domain?

An authentication domain is a name that facilitates the matching of logins with the servers for which they are valid.

Note: This matching is not important when you launch a client, but it is important when you access certain secondary servers such as a third-party DBMS or, in some configurations, a standard workspace server.

Each user ID and password is valid within a specific scope. For example, the user ID and password that you use to log on to your computer at work are probably not the same as the user ID and password that you use at home. It is also common for database servers and web servers to have their own authentication mechanisms, which require yet another, different, user ID and password.

An enterprise application that provides access to many different resources might require that a user have several sets of credentials. Each time a user requests access to a resource, the software must determine which credentials to use to provide access. The software could challenge the user with an interactive prompt for user ID and password, but that quickly becomes an annoyance that interrupts the user experience. The software could randomly try different credentials until it finds a set that works, but authentication attempts can be expensive in terms of performance. In SAS Intelligence Platform, the software attempts to use only the credentials that it expects to be valid for a particular resource or system.

The software's knowledge of which credentials are likely to be valid is based entirely on authentication domain assignments. For this reason, you must correctly assign an authentication domain to each set of resources that uses a particular authentication provider. You must also assign that same authentication domain to any stored credentials that are valid for that provider.

When Do I Need to Add an Authentication Domain?

In the simplest case, all logins and SAS servers are associated with one authentication domain (DefaultAuth). This list describes the most common reasons for using more authentication domains:

- If you use web authentication, you might need a second authentication domain for the logins that contain web-realm user IDs.
- If you have a third-party server (such as a DBMS server) that has its own user registry, you need a separate authentication domain for that server and its logins.
- If both of the following criteria are met, you need a separate authentication domain for the standard workspace server and its logins:
  - The standard workspace server does not share an authentication provider with the metadata server (and cannot be configured to do so).
  - You want to provide seamless individualized access to the standard workspace server.

About Passwords

Managed Passwords

Passwords for a few required accounts (such as the SAS Administrator and the SAS Trusted User) are included in configuration files. If these passwords change, you must also use SAS Deployment Manager to update the configuration. For instructions, see *SAS Intelligence Platform: Security Administration Guide*.

Passwords in Logins

It is usually not necessary to store an external password in the SAS metadata. The main reason to include a password in a login is to provide seamless access to a server that requires credentials that are different from the credentials that users initially submit. The most common example is a deployment that includes a third-party DBMS server that requires a different set of credentials.

Password management for logins is driven by changes that occur in other systems. For example, if you have a personal login for a third-party DBMS, and you change your DBMS password, you must also update the SAS copy of that password.

If credentials are not otherwise available, some applications prompt users for an appropriate user ID and password.

Passwords in Internal Accounts

Internal accounts exist only in the metadata. Each internal account includes a password. By initial policy, internal passwords do not expire.

Passwords in Configuration Files

Passwords for a few required accounts (such as the SAS Administrator and the SAS Trusted User) are included in configuration files. If you need to change these passwords, use SAS Deployment Manager. For instructions, see *SAS Intelligence Platform: Security Administration Guide*.

About External Identities

What is an External Identity?

Logins and internal accounts are involved in the logon process, but external identities are not. An external identity is an optional synchronization key for a user, group, or role. If you use batch processes to coordinate SAS identity information with your primary user registry, you need external identities (such as employee IDs) to facilitate matching. This list explains the circumstances in which a user, group, or role needs an external identity:

- For a user, group, or role that you maintain interactively, no external identity is needed.
- For a user, group, or role that you maintain using batch processes (user import macros), one external identity is needed.

Where do External Identities Come From?

External identities can be added in these ways:

- For a user, group, or role that is created by an import process, an external identity is added as part of that process.

- For any user, group, or role, you can interactively add an external identity on the **External Identities** page.

Requirement: Unique Names and IDs

Within a metadata server, the following uniqueness requirements apply:

- You cannot create a user definition that has the same name as an existing user definition.
- You cannot create a group or role definition that has the same name as an existing group or role definition.
- You cannot assign the same user ID to different users or groups. All of the logins that include a particular user ID must be owned by the same identity. In this situation, the metadata server resolves each user ID to a single identity.

  Note: The exception to this requirement is logins that are associated with outbound authentication domains. These logins are not subject to these constraints.

  - This requirement is case insensitive. For example, you cannot assign a login with a user ID of **`smith`** to one user and a login with a user ID of **`SMITH`** to another user.
  - This requirement applies to the qualified form of the user ID. For example, you can assign a login with a user ID of **`winDEV\brown`** to one user and a login with a user ID of **`winPROD\brown`** to another user.
- If you give a user two logins that contain the same user ID, the logins must be in different authentication domains. Within an authentication domain, each user ID must be unique. For example, if you give Tara O'Toole two logins that both have a user ID of **`tara`**, then you cannot associate both of those logins with the OraAuth authentication domain. As with the previous requirement, this requirement is case insensitive and is applied to the fully qualified form of the user ID.

TIP Avoid using spaces or special characters in the name of a user, group, or role that you create. Not all components support spaces and special characters in identity names.

Chapter 3

Managing Metadata Access

Overview of Metadata Administration 21
Apply an ACT 22
Create an ACT 22
Update an ACT 23
Add an Explicit Grant or Denial 24
About Metadata Administration 24
About Access Management 24
About ACTs 25
About Granularity and Mechanics 26
About Inheritance and Precedence 27
About Use and Enforcement of Each Permission 28
About Permission Conditions 29
About the Permissions Inspector 29
Icons in Access Management 30
How are Denies and Grants Indicated? 30
How are Direct Controls Indicated? 30
What Does a Blank Cell in an ACT Pattern Mean? 31
Permission Origins 32
Introduction to Permission Origins 32
Simple Permission Origins 32
Inherited Permission Origins 33
Best Practices for Permissions 34
Assign Access Controls to Groups 34
Use Folders to Organize Content 34
Centralize Permissions with Access Control Templates 35
Deny Broadly, Grant Selectively (To the Extent Possible) 35

Overview of Metadata Administration

This application supports the access control tasks that are provided by the Authorization module to SAS Environment Manager, including the following:

- applying access control templates (ACTs) to metadata objects
- applying explicit controls to objects

- managing repository-level controls
- maintaining ACTs

In addition, this application provides new access management features, including the following:

- The basic authorization display provides a full grid of applicable permissions and access control participants. This enables you to immediately see the entire picture, instead of having to examine the settings for only one identity at a time. You can immediately see the impact of your access control changes across identities. For example, the impact that an explicit denial for PUBLIC has on all restricted identities that do not have offsetting direct controls is immediately apparent.
- For each effective permission, you can view origins information that identifies the source of the effective grant or denial.
- The permissions inspector, enables you to easily and safely look up effective permissions for any user or group.
- In each ACT's definition, a usage tab lists the objects to which that ACT is directly applied. This helps you identify any gaps in your access control implementation, and helps you anticipate the impact of any changes that you make to the ACT.
- In each object's authorization properties, a summary tab provides a simple list of any direct controls (explicit controls and directly applied ACTs) that have been set on the object. You can rearrange the list to group settings by identity, permission, or type (explicit versus ACT).

Note: This application provides improved access to information. It does not introduce any changes to the underlying security model.

Apply an ACT

1. Open the object to which you are applying the ACT.
2. From the drop-down menu, select **Apply ACT**. On the Apply ACT page, select the check box for the appropriate ACT.

   TIP You should not apply the repository ACT (which is usually named Default ACT) directly to any object. The repository ACT participates through inheritance, serving as an access control parent of last resort.
3. From the drop-down menu, select **ACT: Pattern**. The ACT: Pattern page lists the identities that participate in the pattern of an ACT. Verify that the revised settings are as you expect.
4. Save your changes.

Create an ACT

1. From the **Folders** module, navigate to **SAS Folders** ⇨ **System** ⇨ **Security** ⇨ **Access Control Templates**.
2. Right-click **Access Control Templates** and select **New Access Control Template**.

3. In the New Access Control Template window, enter a name and description for the ACT. Click **OK**.
4. Right-click the new ACT and click **Open**. The new ACT opens.
5. From the drop-down menu, select **ACT: Pattern**. On the **ACT: Pattern** page:
   a. Click ➕.
   b. In the Add Identities window, search for the identities that you want to add. Enter the first few characters of the identity name and click 🔍.
   c. From the list of identities that meet your search criteria, double-click the users and groups that will have explicit settings in the pattern. Click **OK**.
   d. On the **ACT: Pattern** page, click cells and make selections from the lists to define the ACT's pattern.

      Note: For information about the icons and blank cells in an ACT's pattern, see "Icons in Access Management" on page 30.
6. Protect the new ACT. For example, one approach is to add an explicit denial of WriteMetadata for PUBLIC and an offsetting explicit grant of WriteMetadata for SAS Administrators.

   Note: It is important to prevent regular users from modifying or removing an ACT.
7. Save the new ACT.
8. To use the ACT, apply it to one or more objects.

   Note: The applied ACT contributes its pattern of access controls to the object's protections. The object can also have explicit controls and other applied ACTs (as well as inherited settings).
9. If necessary, adjust the ACT's pattern. You can change the pattern of an ACT without modifying the objects to which the pattern is applied.

Update an ACT

CAUTION:
One ACT can protect thousands of objects. Changes that you make to an ACT's pattern affect every object to which that ACT is applied.

Locate an ACT

1. From the **Folders** module, navigate to **SAS Folders** ⇨ **System** ⇨ **Security** ⇨ **Access Control Templates**.

   Note: To locate ACTs that are in custom repositories, your navigation path will vary slightly. For example: **SAS Folders** ⇨ ***custom-repository*** ⇨ **System** ⇨ **Security** ⇨ **Access Control Templates**.
2. Find the ACT that you want to update.

Modify an ACT

3. Right-click the ACT and select **Open**. The ACT opens.
4. From the drop-down menu, select **ACT: Usage**. Examine this page to understand the potential impact of your intended changes.

5. From the drop-down menu, select **ACT: Pattern**. Adjust settings modify the ACT's pattern.

   Note: For information about the icons and blank cells in an ACT's pattern, see "Icons in Access Management" on page 30.

6. Save the ACT.
7. (Optional) Navigate to an object that uses the ACT and verify that the revised settings are as you expect.

Note: To delete an existing ACT, use SAS Management Console. For more information, see *SAS Management Console: Guide to Users and Permissions*.

Add an Explicit Grant or Denial

1. Open the object that you want to protect or make available.
2. From the drop-down menu, select **Authorization**. Locate the user or group that you want to assign an explicit control to. If the user or group is not listed, click ➕ to open the Add Identities window.

   Note: An explicit grant of the ReadMetadata permission is automatically set for each identity that you add.

3. Click a cell and make a selection from the list.

   Note: If the selected identity is unrestricted, all permissions are granted and you cannot make changes.

   Note: When you click outside the cell, ◆ indicates an explicit control is displayed in the cell that you updated.

4. If you changed the access for a group, review the impact on all of the listed identities.

   Note: Controls that you add for a group can affect access for all members of that group. For example, an explicit denial that you add for the PUBLIC group blocks access for all restricted users, unless there are also explicit (or direct ACT) grants. You must offset a broad explicit denial with explicit (or direct ACT) grants for any restricted identities whose access you want to preserve.

5. Save your changes.

TIP It is easy to add explicit grants and denials on each object that you want to protect or make available. However, adding a large number of individual access controls can make access control management cumbersome.

About Metadata Administration

About Access Management

You can use SAS Environment Manager to manage access in the metadata authorization layer. The access control tasks that are provided by SAS Environment Manager include:

- application of access control templates (ACTs) to objects

- maintenance of ACTs
- application of explicit controls to objects
- management of repository-level controls

Over the lifecycle of SAS 9.4, functions will be added to extend SAS Environment Manager's capabilities as a centralized administration application for all SAS products. SAS Environment Manager is not currently a replacement for SAS Management Console, and no functionality has been removed from SAS Management Console.

The following topics provide a brief overview of the metadata authorization model. For a comprehensive discussion, see the *SAS Intelligence Platform: Security Administration Guide*.

Access management determines which objects a user can see and interact with. Permissions that you set on the Authorization page are part of a metadata-based access control system within the SAS Metadata Server.

The SAS metadata authorization layer supplements protections in other layers (such as the operating system, a third-party DBMS, or the SAS Content Server). Protections are cumulative across layers. You cannot perform a task unless you have sufficient access in all layers.

CAUTION:

Do not rely exclusively on the metadata authorization layer to protect data. You must manage physical access (operating system and DBMS permissions) in addition to metadata layer access.

About ACTs

Why Use ACTs?

Use ACTs to avoid having to repeatedly add the same explicit controls for the same identities on multiple objects. When you apply an ACT to an object, the pattern settings in an ACT are added to the direct controls of an object.

TIP Settings in the pattern of an ACT affect access to all of the objects to which the ACT is applied. Settings on the Authorization page for an ACT affect who can access that ACT.

Why Create Custom ACTs?

Several predefined ACTs are provided. To further centralize access management, create an ACT for each access pattern that you use repeatedly. Here are some common patterns and tips:

- It is often useful to create ACTs to manage Read access for different business units.
- It is often useful to create an ACT that manages Write access for a functional group that includes users from multiple business units.
- You do not have to capture all of an object's protections in one ACT. You can use combinations of ACTs, explicit controls, and inherited settings to define access to an object.

About Granularity and Mechanics

Repository-Level Controls

Repository-level controls function as a gateway. Participating users need ReadMetadata and WriteMetadata permissions for the foundation repository.

Repository-level controls also serve as a parent of last resort, defining access to any objects that do not have more specific settings. Repository-level controls are displayed on the ACT: Pattern page of the repository access control template (Default ACT).

Why Adjust the Repository-Level Controls?

CAUTION:

Altering the repository-level controls for service identities can prevent necessary access. We recommend that you do not change these settings.

Here are some key points about working with repository-level controls for a foundation repository:

- If you want some or all users to have default Read access to all data, grant the Read permission at the repository level.
- If you want to experiment with changing repository-level access, we recommend that you create a new ACT and designate that ACT as the repository ACT (instead of modifying the original repository ACT).
- All users need ReadMetadata and WriteMetadata access to the foundation repository. It is appropriate for the SASUSERS group to be granted these permissions in the pattern of the repository ACT.

Which ACT is the Repository ACT?

If your site has multiple metadata repositories, you have multiple repository ACTs. Each repository has its own repository ACT, which is usually named Default ACT.

As an alternative to opening each ACT to determine which repository it belongs to, navigate to the ACT from within the **Folders** module.

- ACTs for the foundation repository are located in the **SAS Folders** ⇨ **System** ⇨ **Security** ⇨ **Access Control Templates** folder.
- ACTs for a custom repository are located in the **SAS Folders** ⇨ ***custom-repository*** ⇨ **System** ⇨ **Security** ⇨ **Access Control Templates** folder.

Note: The repository ACT indictor is located at the bottom of the ACT: Usage page for *access-control-template*.

Object-Level Controls

Object-level controls manage access to a specific object such as a report, an information map, a stored process, a table, a cube, or a folder. You can define object-level controls individually (as explicit controls) and in patterns (as directly applied access control templates).

Fine-Grained Controls

Fine-grained controls affect access to subsets of data within an object. To establish fine-grained controls, you define permission conditions that constrain access to rows within a table or members within an OLAP dimension.

Feature-Level Controls

Some applications use roles to limit access to functionality. These applications check the roles of each user in order to determine which menu items and features to display for that user. Roles management is part of user administration.

About Inheritance and Precedence

Two Relationship Networks

Permissions are conveyed across two distinct relationship networks—a resource network and an identity network. Permissions that are set directly on an object always have priority over permissions that are set on the parent of an object. For example, when access to a report is evaluated, a denial that is set on the report (and explicitly assigned to the PUBLIC group) overrides a grant that is set on the parent folder of the report (even if the grant is explicitly assigned to you).

The Resource Relationships Network

Permissions that you set on one object can affect access to many other objects. For example, a report inherits permissions from the folder in which the report is located. This relationship network consists primarily of the SAS Folders tree. This list highlights some exceptions:

- The root folder is not the ultimate parent. This folder inherits from the repository (through the permission pattern of the repository access control template (ACT)).
- The root folder is not a universal parent. Some system resources (such as application servers, identities, and ACTs) have the repository as their immediate and only parent.
- Inheritance within a table or cube follows the data structure. For example, table columns and cube hierarchies do not have a folder as their immediate parent. Instead, a column inherits from its parent table and a hierarchy inherits from its parent cube.
- In unusual circumstances, it is possible for an object to have more than one immediate parent. If there is a tie in this network (for example, if there are no settings on an object, the object has two immediate parents, and one parent provides a grant while the other parent provides a denial), the outcome is a grant. If there are no direct controls, a grant from any inheritance path is sufficient to provide access.
- In general, specialized folders (such as search folders, favorites folders, and virtual folders) do not convey permissions to the objects that they contain. An exception is that a favorites folder does convey permissions to any child favorites folders (favorites groups) that it contains.

The Identity Relationships Network

Permissions that you assign to one group can affect access for many other identities. For example, if you grant a group access to an OLAP cube, that grant applies to all users who are members of the group. This relationship network is governed by a precedence order that starts with a primary identity, can incorporate multiple levels of group memberships, and ends with implicit memberships in SASUSERS and then PUBLIC. If there is a tie in this network (for example, if you directly assign a user to two groups and give one group an explicit grant and another group an explicit denial), the outcome is a denial.

About Use and Enforcement of Each Permission

Table 3.1 *Permission Reference*

Permission (Abbreviation)	Actions Affected and Limitations on Enforcement
ReadMetadata (RM)	View an object. For example, to see an information map, you need RM for that information map. To see (or traverse) a folder, you need RM for that folder.
WriteMetadata (WM)	Edit, delete, change permissions for, or rename an object. For example, to edit a report, you need WM for the report. To delete a report, you need WM for the report (and WMM for the parent folder of the report). WM can also affect the ability to create associations. For example, you need WM on an application server in order to associate a library with that server. WM affects the ability to create objects in certain containers. For example, to add an object anywhere in a repository, you need WM at the repository level. For folders, adding and deleting child objects is controlled by WMM, not WM.
WriteMemberMetadata (WMM)	Add an object to a folder or delete an object from a folder. For example, to save a report to a folder, you need WMM for the folder. To remove a report from a folder, you need WMM for the folder (and WM for the report). To enable someone to interact with the contents of a folder, but with not the folder itself, grant WMM and deny WM.*
CheckInMetadata (CM)	Check in and check out objects in a change-managed area. Change management is an optional feature that is supported by only SAS Data Integration Studio.**
Administer (A)	Monitor, stop, pause, resume, refresh, or quiesce a server or spawner. For the metadata server, the ability to perform tasks other than monitoring is managed by the **Metadata Server: Operation** role (not by this permission).
Read (R)	Read data. For example, you need RM for a cube in order to see the cube, and you need R for the cube in order to run a query against it. Enforced for OLAP data, information maps, data that is accessed through the metadata LIBNAME engine, and dashboard objects.
Create (C)	Add data. For example, on a table, C controls adding rows to the table. Enforced for data that is accessed through the metadata LIBNAME engine.
Write (W)	Update data. For example, on a table, W controls updating the rows in the table. Enforced for data that is accessed through the metadata LIBNAME engine, for publishing channels, and for dashboard objects.
Delete (D)	Delete data. For example, D on a library controls the deletion of tables from the library. Enforced for data that is accessed through the metadata LIBNAME engine and for dashboard objects.

Permission (Abbreviation)	Actions Affected and Limitations on Enforcement
ManageMemberMetadata (MMM)	Change the membership of the Group and Role. Cannot change security or other account attributes
ManageCredentialsMetadata (MCM)	Manage accounts and trusted logins of User and Group. Cannot change security or other account attributes.

* A folder's WMM settings mirror its WM settings unless the folder has a direct control for WMM. A grant (or deny) of WMM on a folder becomes an inherited grant (or deny) of WM on the objects and subfolders within that folder. WMM is not inherited from one folder to another. WMM is not applicable to specialized folders (such as virtual folders, favorites folders, or search folders).

** In any change-managed areas of a foundation repository, change-managed users should have CM (instead of WM and WMM).

Note: For information about the Insert, Update, Select, Create Table, Drop Table, and Alter Table permissions, and an additional use of the Delete permission, see the *SAS Guide to Metadata-Bound Libraries*.

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

About Permission Conditions

What is a Permission Condition?

A permission condition limits an explicit grant of the Read permission so that different users access different subsets of data.

About the Permissions Inspector

The permissions inspector enables you to easily and safely look up effective permissions for any user or group.

To launch the inspector, click .

The inspector offers the following features:

- The inspector shows the effective permissions that a selected identity has for the specified object. Permissions information is displayed after you look up and select an identity.
- The contents are also updated when you look up and select a different identity (in the text box within the inspector).
- The inspector uses the same icons and indicators that are used on the Authorization page.
- You can view permission origin information by clicking the grant or deny icon.

Here are some tips for using the inspector:

- The inspector is always read-only. To set permissions, open the target object and use its authorization tabs.
- To select an identity, enter a user or group name in the text box. The search is against display name (or, for an identity that does not have a display name, name). The search uses the "contains" criteria — so that you can provide any part of the name.
- Conditional grants are indicated in the inspector, but you cannot access the associated permission conditions from the inspector. Use the Authorization page to

view or update a permission condition. Permission conditions can be applied to both LASR tables and secured tables.

- One inspector window can be open for each object.
- The inspector does not reflect unsaved changes.

Icons in Access Management

How are Denies and Grants Indicated?

Table 3.2 *Denies and Grants*

Icon	Meaning
⊘	Deny
⊙	Grant

How are Direct Controls Indicated?

The main displays of effective permissions on the Authorization page use the following icons to provide immediate information about the source of each setting.

Table 3.3 *Direct Access Controls*

Icon	Term	Meaning
◆	Direct control: Explicit	The direct access control is set on the current object and specifically assigned to the selected identity.
◘	Direct control: ACT	The direct access control comes from an applied access control template (ACT) whose pattern specifically assigns the grant or deny to the selected identity.
(none)	Indirect setting	The setting comes from someone else (a group that has a direct control), somewhere else (a parent object or the repository ACT), or special status (such as unrestricted). For the WriteMemberMetadata permission, indirect means that the setting mirrors the WriteMetadata setting.

TIP The explicit and ACT indicator icons correspond to the white and green colors on the **Authorization** window in SAS Environment Manager. As in SAS Environment Manager, if both an explicit control and an applied ACT setting are present, only the explicit indicator is displayed.

Table 3.4 *Icon Combinations in the Main Authorization Displays*

Icon	Meaning
⊘◆	Deny from an explicit control
⊘□	Deny from an applied ACT
⊘	Deny from an indirect source (such as a parent group or parent object)
⊙◆	Grant from an explicit control
⊙□	Grant from an applied ACT
⊙	Grant from an indirect source (such as a parent group or parent object)

TIP For additional details about the source of a setting, use the permission origins feature.

What Does a Blank Cell in an ACT Pattern Mean?

The display of an ACT's pattern is limited as follows:

- An ACT's pattern includes only those identities that have pattern settings. For this reason, the table on an ACT's **Authorization** ⇨ **Basic** tab usually includes only a few groups. Not all users and groups are listed.
- An ACT's pattern consists of only those settings that are explicitly defined in the pattern. For this reason, the table on an ACT's **Authorization** ⇨ **Basic** tab usually has grants or denies in only a few cells. The other cells are blank.

  Note: This differs from the display in SAS Environment Manager, where the net effect of the pattern is displayed along with the pattern settings.

For each blank cell and each unlisted identity, the net effect of the pattern is determined by the closest pattern setting. Each identity's group memberships determine which setting is closest. The precedence order is as follows:

1. The identity's direct group membership have the highest precedence.
2. The identity's nested group memberships are next, with each successive level of nesting having lower precedence than the preceding level. Nested memberships are a consideration only if the identity is a member of a group that is in turn a member of another group.
3. The identity's automatic membership in the SASUSERS implicit group is next, unless the identity is a user who is not properly registered in the metadata. This group includes all registered users. For example, most users get their repository-level access through grants to SASUSERS in the default ACT's pattern.
4. The identity's automatic membership in the PUBLIC implicit group is last. PUBLIC is a superset of SASUSERS. PUBLIC includes everyone that can connect to the metadata server, regardless of whether they are registered users. Because PUBLIC is the broadest group, denies are usually assigned to it.

If an identity has conflicting pattern settings at the same level of precedence, the net effect of those settings is a deny. If there are no pattern settings that are relevant for an identity, the ACT has no effect on that identity.

Permission Origins

Introduction to Permission Origins

The permission origins feature identifies the source of each effective permission. Permission origins answers the question: Why is this identity granted (or denied) this permission?

In the origins answer, only the controlling (winning, highest precedence) access control is shown. If there are multiple tied winning controls, they are all shown. Other, lower precedence controls are not shown in the origins answer.

Simple Permission Origins

The following table provides simple examples of permission origins answers. In each example, we are interested in why UserA has an effective grant on FolderA. In each example, UserA is a direct member of both GroupA and GroupB. Each row in the table is for a different (independent) permissions scenario. In the table, the first column depicts the contents of the Origins window. The second column interprets the information.

Table 3.5 *Origins: Simple Examples*

Origins Information	Source of UserA's Effective Grant on FolderA
UserA [Explicit]	On FolderA, an explicit grant for UserA
GroupA [Explicit]	On FolderA, an explicit grant for GroupA
GroupA [Explicit] GroupB [Explicit]	On FolderA, explicit grants for GroupA and GroupB *Note:* Two settings are shown because they are tied and they both win (UserA is a direct member of GroupA and GroupB).
GroupA [ACT: GroupARead]	On FolderA, an ACT pattern grant for GroupA (from a directly applied ACT)
SASUSERS [ACT: GenRead]	On FolderA, an ACT pattern grant for SASUSERS (from a directly applied ACT)
GroupA [ACT: GroupARead] GroupB [ACT: GroupBRead]	On FolderA, ACT pattern grants for GroupA and GroupB (from two different directly applied ACTs). *Note:* Two settings are shown because they are tied and they both win (UserA is a direct member of GroupA and GroupB).

Origins Information	Source of UserA's Effective Grant on FolderA
GroupA [ACT: GroupABRead] GroupB [ACT: GroupABRead]	On FolderA, ACT pattern grants for GroupA and GroupB (from the same directly applied ACT). *Note:* Two settings are shown because they are tied and they both win (UserA is a direct member of GroupA and GroupB).
UserA is unrestricted.	UserA's status as an unrestricted user (someone who is unrestricted is always granted all permissions)
This setting mirrors the WriteMetadata setting.	The WriteMetadata setting on FolderA. To investigate further, examine the origins for that setting.

Inherited Permission Origins

In many cases, the controlling setting is not on the current object. Instead, the controlling setting is defined on a parent object and inherited by the current object.

The following table provides examples in which the controlling setting comes from a parent object. Because the source of the effective permission is a parent object, the answer must identify which parent object has the controlling setting. For this reason, the origins answers in the following examples identify both a particular parent object (the object that has the controlling setting) and the controlling setting itself.

In each example, we are interested in why UserA has an effective grant on FolderA. In each example, UserA is a direct member of both GroupA and GroupB. Each row in the table is for a different (independent) permissions scenario. In the table, the first column depicts the contents of the Origins window. The second column interprets the information.

Table 3.6 *Origins: Inheritance Examples*

Origins Information	Source of UserA's Effective Grant on FolderA
ParentFolderA UserA [Explicit]	On ParentFolderA, an explicit grant for UserA
ParentFolderA GroupA [Explicit]	On ParentFolderA, an explicit grant for GroupA
ParentFolderA GroupA [Explicit] GroupB [Explicit]	On ParentFolderA, explicit grants for GroupA and GroupB
ParentFolderA GroupA [ACT: GroupARead]	On ParentFolderA, an ACT pattern grant for GroupA (from a directly applied ACT)

Origins Information	Source of UserA's Effective Grant on FolderA
GreatGrandParentFolderA SASUSERS [ACT: GenRead]	On GreatGrandParentFolderA, an ACT pattern grant for SASUSERS (from a directly applied ACT)
ParentFolderA GroupA [ACT: GroupARead] GroupB [ACT: GroupBRead]	On ParentFolderA, ACT pattern grants for GroupA and GroupB (from two different directly applied ACTs)
GrandParentFolderA GroupA [ACT: GroupABRead] GroupB [ACT: GroupABRead]	On GrandParentFolderA, ACT pattern grants for GroupA and GroupB (from the same directly applied ACT).
SAS Folders SASUSERS [Explicit] Default ACT [CustomRepositoryA] SASUSERS [Pattern]	On the **SAS Folders** node, an explicit grant for SASUSERS. Also, in CustomRepositoryA's default ACT, a pattern grant for UserA. *Note:* In this example, FolderA is within a custom repository, so it inherits from both the **SAS Folders** node and the custom repository's default ACT pattern. Two settings are shown because they are tied and they both win.

Best Practices for Permissions

Assign Access Controls to Groups

You can simplify access control by assigning permissions to groups rather than to individual users. Here are some examples:

- To allow only unrestricted users to access an object, assign denials to the PUBLIC group.
- To enable only registered users to access an object, assign denials to the PUBLIC group and grants to the SASUSERS group.
- To enable only information developers and unrestricted users to access an object, create a group for the developers. On the object, assign denials to the PUBLIC group and assign grants to the developers group.

Use Folders to Organize Content

You can simplify access control by creating a folder structure that reflects the access distinctions that you want to make. Instead of adding access controls to each individual object, add access controls to parent folders. The objects in a folder inherit the effective permissions for a folder.

TIP To protect the folder structure, do not grant WriteMetadata permission on a folder to someone if granting them WriteMemberMetadata permission is sufficient.

Centralize Permissions with Access Control Templates

You can simplify access control by using access control templates (ACTs). An ACT is a reusable pattern of settings that you can apply to multiple objects. Each ACT consists of the following elements:

- a list of users and groups
- for each identity and permission, an indication of whether the pattern of an ACT provides a grant, a denial, or no pattern setting

Deny Broadly, Grant Selectively (To the Extent Possible)

Assign denials to the broadest group (PUBLIC) and then add offsetting grants for users or groups whose access you want to preserve. Deny access at the highest point of control and then grant access back on specific containers or objects. The following constraints apply:

- The highest point of control is the repository-level settings for the foundation repository. The security model requires that participating users have ReadMetadata and WriteMetadata access at this level. Therefore, broadly denying access in the permission pattern of the foundation repository's Default ACT is not a workable approach. Instead, use the next point of control, which is the top of the **SAS Folders** tree.
- In order to navigate within the folder tree, users need a clear path of grants of ReadMetadata to the objects that they use. For the ReadMetadata permission, setting denials on folders at a high level is not a workable approach.

  Note: When users access objects by searching (instead of by navigating through the folder tree), a clear path of grants of ReadMetadata is not required.

Chapter 4

Managing Servers

Using the Servers Module 37
Create a New Server Definition 37
Filtering the Server List 38

Using the Servers Module

The Servers module enables you to manage server definitions in metadata. For the current release, you can browse any type of server that has been defined in SAS metadata. You can create and edit definitions for SAS LASR Analytic Servers.

To open the Servers module, click the **Side Menu** icon in the SAS Environment Manager banner and select **Servers** from the side menu. The Servers view displays a tree view of all server definitions in the SAS Metadata Server. Click the **Filter** icon to filter the view by server type.

You can expand or collapse individual entries or all entries (using the controls above the list). Double-click any object in the list to view its properties.

Create a New Server Definition

Follow these steps to create a new server definition in metadata:

1. In the Servers module, click the **New Server** icon to display the New Server window.
2. Specify the appropriate information in the New Server window. The window contains only the minimum fields required to define the server. When you select the server type, the **Options** section of the window automatically expands to display any required options for the selected server type. SAS LASR Analytic Server is currently the only available server type.
3. Click **Save** to create the definition. The new server definition automatically opens to enable you to specify any non-required options..
4. To specify any non-required options, click the property page title to see the list of all property pages for the server. Specify any additional options and save the definition.

Note: If you are using Microsoft Internet Explorer 11, you might need to click the **Refresh** icon in order to see the new server definitions in the list of servers.

After you use the Servers module to create the SAS LASR Analytic Server definition, you must complete these additional steps in order for SAS Visual Analytics to be able to use the server:

1. Use the **Authorization** page to set the metadata permissions for users of the server. They must have ReadMetadata, WriteMetadata, and Administer permissions.
2. Use the **Authorization** page to add the Visual Analytics Data Administrators group to the identity list for the server. You must search for this group in order to add it.
3. Grant the Administer permission for the Visual Analytics Data Administrators group.

See SAS 9.4 Intelligence Platform: Data Administration Guide for more information.

Filtering the Server List

If your system has many servers defined in metadata, it can be difficult to locate a specific server definition or all definitions for a particular server type. You can filter the list of servers to make this task easier.

Click the filter icon on the toolbar and select the type of server you want to display. The server list then displays only servers of the selected type. To display all the defined servers, select **All** as the server type.

Chapter 5

Managing Libraries

Using the Libraries Module 39

Creating a New Library Definition 39

Registering a Table 40
Register a Table from the List of Libraries 40
Register a Table from a Library's Properties Page 41
Register a Table from the Folders Module 41

Change Assigned SAS Servers 41

Using the Libraries Module

The Libraries module enables you to manage SAS library definitions in metadata. For the current release, you can browse any type of library that has been defined in SAS metadata. You can create and edit definitions for Base SAS libraries and SAS LASR Analytic Server libraries.

To open the Libraries module, click the **Side Menu** icon in the SAS Environment Manager banner and select **Libraries** from the side menu. The Libraries view displays a table of all library definitions in the SAS Metadata Server. Click the Filter icon to filter the table by library type. You can also search the table, sort the table by a selected column, and choose which columns appear in the table. Double-click on a table entry to open the library definition.

To delete a library definition, either select the entry in the table and click the **Delete** icon or right-click the entry and select **Delete** from the pop-up menu. A dialog box appears for you to confirm the deletion. If any tables have been registered with the library, you must confirm the deletion again. If you delete a library to which tables have been registered, the table objects are deleted as well.

Creating a New Library Definition

Follow these steps to create a new library definition:

1. In the Libraries module, click the **New Library** icon to display the New Library window.

2. Specify the requested information in the New Library window. The window contains only the fields that are required to define the library. Required fields are identified by a red asterisk next to the field name.

   When you select the library type, the **Options** section of the window automatically expands to display any required options for the selected library type.

   Note: For the current release, **SAS LASR Analytic Server Library** and **SAS BASE Library** are the only two available values for the **Type** field.

3. Click **OK** to create the definition. After the definition is created, it automatically opens to enable you to specify any non-required options.

4. To specify non-required options, click the property page title to see the list of all property pages for the library. Specify any additional options and save the definition.

Registering a Table

Before you can register a table in a library, you must be logged on to SAS Environment Manager using an ID that can register libraries. You must also assign the library definition to a SAS Application Server. If you are registering tables for a SAS LASR library, previously registered tables might be included in the list of available tables. If you attempt to register these tables again, the registration will fail, although the table is still registered with the library. You can register a table with a library in three ways:

Register a Table from the List of Libraries

1. In the Libraries module, right-click the library that should contain the table. Select **Register Table** from the pop-up menu. The Register Tables dialog box appears.

2. The default location for data tables is **My Folder** for the current user. Click **Browse** to select another location.

   The location is where the metadata for the table is stored. The tables that are being registered are retrieved from the location specified in the **Location** field in the library definition.

3. Select tables to register from the **Available tables** list.

4. Select these options as needed:

   Enable case-sensitive DBMS object names
   specifies whether case-sensitive names for tables and columns are supported in the metadata that you are about to generate for tables in the current library. If the check box is deselected, no support is provided. If the check box is selected, support is provided.

   Enable special characters within DBMS object names
   specifies whether special characters in names for tables and columns are supported in the metadata that you are about to generate for tables in the current library. If the check box is deselected, no support is provided. If the check box is selected, support is provided.

5. Click **OK** to register the selected tables.

Register a Table from a Library's Properties Page

1. From the Libraries module, double-click an entry to open the Properties page for the library.
2. Click the page title (**Basic Properties**) and select **Tables** from the list of property pages. The list of tables registered to the library appears.
3. Click the **Register Tables** icon ➕ to register a table. The Register Tables dialog box appears.
4. The default location for data tables is **My Folder** for the current user. Click **Browse** to select another location.
5. Select tables to register from the **Available tables** list.
6. Select these options as needed:
   - **Enable case-sensitive DBMS object names**
   - **Enable special characters within DBMS object names**
7. Click **OK** to register the selected tables.

Register a Table from the Folders Module

1. Use the side menu to open the **Folders** module.
2. Navigate in the **Folders** tree until you locate the library that should contain the table.
3. Right-click the library entry and select **Register Tables** from the menu.
4. The default location for data tables is **My Folder** for the current user. Click **Browse** to select another location.
5. Select tables to register from the **Available tables** list.
6. Select these options as needed:
   - **Enable case-sensitive DBMS object names**
   - **Enable special characters within DBMS object names**
7. Click **OK** to register the selected tables.

Change Assigned SAS Servers

Each library can be assigned to one or more defined SAS servers. Assigning a library to a server means that the server is aware of the library and has access to the library.

Follow these steps to change the SAS servers to which the library is assigned:

1. In the Libraries module, double-click a library definition to open it.
2. In the menu bar, click **Basic Properties** (the title of the default property page) and select **Assigned SAS Servers** from the menu.
3. The property page lists all of the defined SAS servers. Select the check boxes for the servers to which the library should be assigned.

Chapter 6

Managing Backups

Overview of SAS Backup Manager **43**
Introduction 43
What is Backed Up? 44
What is Not Backed Up? 45
Default Backup Schedule 45
Default Backup Location 45
About Central Vault Locations 45
Backup Alert Notifications 46
Accessing Logs 46
Where to Obtain More Information 46

Viewing Backup and Recovery History **47**
View the History List 47
Sort the History List 48
Prevent Column Scrolling 48
View Backup or Recovery Operation Details 48
View Backup or Recovery Sources 49

Running an Unscheduled Backup **50**

Viewing and Modifying the Backup Policy **51**
View the Backup Policy 51
Update the Backup Policy 51
View Configuration Information for Backup Sources 53

Viewing and Modifying the Backup Schedule **53**
View the Backup Schedule 53
Best Practices for Scheduling Backups 53
Enable or Disable Scheduled Backups 54
Update the Backup Schedule 54

Overview of SAS Backup Manager

Introduction

SAS Backup Manager is a user interface that enables you to schedule, configure, monitor, and perform integrated backups of your SAS content across multiple tiers and machines. This interface, which is new with the third maintenance release of SAS 9.4, enables you to perform most of the functions of the Deployment Backup and Recovery

tool. In previous SAS releases, these functions were available only through batch commands.

To open SAS Backup Manager, click the **Side Menu** icon ☰ in the SAS Environment Manager banner and select **SAS Backup Manager** from the side menu. SAS Backup Manager takes several minutes to discover assets in your deployment that are available for backup. You can then access the following pages:

Table 6.1 *SAS Backup Manager Functionality*

Page	Description
History	• view a list of backups or recoveries that have been run, are currently running, or are waiting to run • view details for a particular backup or recovery operation • view information about sources for a particular backup or recovery • view a tree diagram of sources for a particular backup or recovery • start an immediate backup
Policy	• view a tree diagram of currently configured backup sources • view details of the current backup policy • update the backup policy • view configuration information for backup sources
Schedule	• view the current backup schedule • modify the backup schedule • enable or disable scheduled backups

What is Backed Up?

The Deployment Backup and Recovery tool, which is the underlying software for SAS Backup Manager, backs up the following components of your deployment:

- the metadata server, including all registered metadata repositories, the repository manager, and the server's configuration directory.
- the contents of the `Data` directories, `SASEnvironment` directories, and server configuration directories for each server on the SAS server tier. (If symbolic links in these directories point to other locations, the referenced locations are not backed up.)
- the SAS Content Server repository.
- the databases that are managed by the SAS Web Infrastructure Platform Data Server. By default, all of the databases are backed up.
- additional directories under ***`SAS-configuration-directory`***`/Levn` as specified by the administrator.

For backups to work, the SAS Deployment Agent must be running on each middle-tier and server-tier host machine. The Deployment Backup and Recovery tool connects with the agent and automatically discovers the tiers in your deployment and their installed components. New components in your deployment are detected automatically and added

to the backup. For example, it detects new instances of the SAS Web Infrastructure Data Server and new databases that are managed by the server.

TIP If you need to exclude specific tiers, servers, databases, directories, or files from the backup, you can do so by using the command sas-update-backup-config.

What is Not Backed Up?

The Deployment Backup and Recovery Tool has the following limitations:

- Host machines on which the SAS Deployment Agent is not installed are automatically excluded from backups.
- The tool backs up only SAS content and configuration information. It does not back up your SAS software.
- If you are using a third-party vendor database (instead of the SAS Web Infrastructure Platform Data Server) for the SharedServices database, the Deployment Backup and Recovery Tool cannot back it up.
- The tool does not back up the entire contents of your SAS configuration directories. It backs up only **`Data`** directories, the **`SASEnvironment`** directories, and the server configuration directories for each server on the SAS server tier.

  TIP To back up additional subdirectories directories under ***`SAS-configuration-directory`***`/Levn`, add them with the command sas-update-backup-config.

Default Backup Schedule

By default, the Deployment Backup and Recovery tool performs a backup automatically each Sunday at 1:00 a.m. and retains backup files for a period of 30 days. To change the schedule, you can use either SAS Backup Manager or sas-set-backup-schedule. To change the retention period, you can use either SAS Backup Manager or sas-update-backup-config.

Default Backup Location

By default, backup files are stored locally on the same machine where the backed up component is located. All components are backed up to the following path on their respective host machines: ***`SAS-configuration-directory`***`/Lev1/Backup/Vault`. This directory is created on each machine the first time a backup is executed.

For metadata server backups, the tool uses the backup files that are created by the metadata server backup utility. The tool copies these files to ***`SAS-configuration-directory`***`/Lev1/Backup/Vault` on the metadata server machine. If metadata server clustering is configured, the files are copied to the initially configured metadata server.

About Central Vault Locations

You can also specify a central, network-accessible vault location to store backups. If a central vault is specified, backups are automatically copied from the various host machines to the central vault following each backup operation.

TIP A central vault location is highly recommended to avoid the loss of backup files in the event that a host machine fails.

If you have a homogeneous operating system environment, you can use either SAS Backup Manager or the sas-update-backup-config command to specify a central vault location. A homogeneous environment is one in which all of the host machines that are included in the backup are in the same operating system family. For example, Solaris and HP-UX machines are both considered to be in the UNIX operating system family.

In heterogeneous (mixed) operating system environments, you must use manual steps to specify a central vault. See "Specifying a Central Vault Location in a Heterogeneous Operating System Environment" in the *SAS Intelligence Platform: System Administration Guide*.

Backup Alert Notifications

An alert email is generated if a backup or recovery is unsuccessful. By default, the email is sent to the system administrator email address that was specified in the SAS Deployment Wizard. You can use either SAS Backup Manager or the sas-update-backup-config command to specify different email addresses.

Accessing Logs

For details about warnings and errors, you can view generated log messages in the following log files:

- The following log file is created on the middle-tier machine:

  ***SAS-configuration-directory*/Lev1/Web/Logs/SASServer1_1/SASDeploymentBackup9.4.log**

  By default, this log reports only errors and warnings. If you want to set different logging levels, you can do so by editing SASDeploymentBackup-log4j.xml, which is located in ***SAS-configuration-directory*/Lev1/Web/Common/LogConfig/**.

- For backup, recovery, and purge operations, log files are created in the directories where local backups are stored. The default location is as follows:

  SAS-configuration-directory*/Lev1/Backup/Logs/backup-ID*

- The following log file might be useful to obtain information about server-side activity:

  ***SAS-configuration-directory*/Lev1/Backup/backupserver.log**

Where to Obtain More Information

See "Using the Deployment Backup and Recovery Tool" in the *SAS Intelligence Platform: System Administration Guide* for more information about backups and recoveries, including:

- an overview of tasks for configuring your backup policy and schedule
- instructions and syntax for the command-line backup tools
- how to manually define a central vault location if you have a heterogeneous (mixed) operating system environment
- how to specify a backup user (to access a central vault location or a clustered metadata server in an environment that includes one or more Windows hosts)

- how to configure settings to support host machines that use a forward proxy to communicate with the SAS Web Server
- how to perform a recovery

The *SAS Intelligence Platform: System Administration Guide* is available at http://support.sas.com/94administration.

Viewing Backup and Recovery History

View the History List

To view a list of backups or recoveries that have been run, are currently running, or are waiting to run:

1. Select **History** from the drop-down menu.
2. At the top of the History page, select **Backup details** or **Recovery details**.

The list includes all backups or recoveries that are recorded in backup history, including backups that have been purged due to the retention policy. It also includes backups or recoveries that are currently running or that are waiting to run. By default, they are listed in descending order by Backup ID or Recovery ID. The following information is listed for each operation:

Status
contains one of the following icons, indicating the status of the backup or recovery operation:

Table 6.2 *Status Icons Used in SAS Backup Manager*

Icon	Meaning
	The backup or recovery has not yet started.
	The backup or recovery is currently running (in progress).
	The backup or recovery completed without errors or warnings.
	The backup or recovery completed with warnings.
	The backup or recovery completed with errors.
	The backup or recovery encountered a fatal error.
	The backup or recovery was canceled.
	The status of the backup or recovery cannot be determined.

TIP Refresh your browser to see the latest status.

Backup ID or **Recovery ID**
the unique identifier of the backup or recovery, based on the date and time that the backup or recovery started (for example, 2015-02-01T03_13_01). For backups, the ID is also the name of corresponding backup directory.

User ID
the user ID of the user that ran the backup or recovery.

Size
the total size of the files that were backed up. This column does not appear when you display recoveries.

Start Time
the date and time that the backup or recovery started running.

TIP You can also use the command sas-list-backups to display a list of backups or recoveries. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

Sort the History List

By default, the backups or recoveries are listed in descending order by Backup ID or Recovery ID. To sort the list in a different order:

1. Right-click a column heading, and select **Sort** ⇨ **Sort (ascending)** or **Sort** ⇨ **Sort (descending)**.
2. To specify a secondary sort field, right-click a different column heading and select **Sort** ⇨ **Add to sort (ascending)** or **Sort** ⇨ **Add to sort (descending)**.

   Repeat this step to add more sort fields as needed.

Prevent Column Scrolling

You can prevent one or more columns from scrolling horizontally. Right-click the heading of the column that you want to remain stationary, and select **Freeze**. The column that you selected, as well as all of the columns to the left of the selected column, will not scroll when you move the scroll bar to the right or to the left.

To allow scrolling, right-click the column heading and select **Unfreeze**.

View Backup or Recovery Operation Details

To view details for a particular backup or recovery operation, select the operation by clicking any of its columns. The following details appear in the right panel:

- the backup ID or recovery ID.
- the backup name, if one was assigned.
- the status of the backup or recovery. See Table 6.2 on page 47.
- the total size of the files that were backed up. This information does not appear for recoveries.
- any comments that were specified when the backup or recovery was run.
- the user ID of the user that ran the backup or recovery.

- the start and end date and time for the backup or recovery.

TIP You can also use the command sas-display-backup to display details about a particular backup. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

View Backup or Recovery Sources

View the List of Sources

The sources for the currently selected backup or recovery are listed in the right pane, below the operation details. If you are viewing details for a recovery, only the sources that were recovered are listed.

The status icon next to each source indicates the status of its backup or recovery. For explanations, see Table 6.2 on page 47.

By default, the backup sources include the following:

Metadata Server
the SAS Metadata Server, including all registered metadata repositories, the repository manager, and the server's configuration directory.

Content Server
the SAS Content Server repository.

Config Directories
the contents of the `Data` directories, `SASEnvironment` directories, and server configuration directories for each server on the SAS server tier. (If symbolic links in these directories point to other locations, the referenced locations are not backed up.)

Database
the databases that are managed by the SAS Web Infrastructure Platform Data Server. By default, all of the databases are backed up. You can modify the backup configuration so that only selected databases are backed up.

Note: If you are using a third-party vendor database (instead of the SAS Web Infrastructure Platform Data Server) for the SharedServices database, the Deployment Backup and Recovery Tool cannot back it up.

Note: **Custom** might also be listed. This means additional directories under ***SAS-configuration-directoryn*/Levn**, as specified by the administrator, were backed up or recovered.

View Details about a Source

To view details about a particular backup or recovery source, click the Collapsed arrow ▶ to the left of the source name. The following details are displayed:

- the host name of the machine where the source is located
- the status of the source's backup or recovery
- the directory location of the source's local backup files on the host machine
- the total size of the backup files for this source
- the directory location of the source's configuration files
- the operating system of the source's host machine

To collapse the source so that the details are not displayed, click the Expanded arrow ◢ to the left of the source name.

View a Diagram of Backup or Recovery Sources

To view a tree diagram of the sources for a particular backup or recovery, click **View Diagram** at the bottom right of the History page. The diagram includes the following:

- The root node specifies the ID of the backup or recovery, which is based on the date and time that the backup or recovery started (for example, 2015-02-01T03_13_01). For backups, the ID is also the name of corresponding backup directory.
- Under the root node, a child node is displayed for each backup source that was included in the backup or recovery.
- Under each source node, a child node is displayed for each host machine for that source.

In the diagram, you can do the following:

- Zoom in or out by clicking the diagram to select it and then pressing **Ctrl** while scrolling the mouse wheel.
- If parts of the diagram are not visible, drag the entire diagram right, left, upward, or downward.
- Hold the mouse pointer on a node to see the size of the files that were baked up or recovered. (For recoveries, size information is not available on the root node.)
- View the status icons on each node to determine the status of the backup or recovery. For explanations, see Table 6.2 on page 47.
- Click a node to collapse its child nodes.
- Click the node again to expand it so that its child nodes reappear.
- Click **Close** to close the diagram.

Running an Unscheduled Backup

To start an immediate backup:

1. Select **History** from the drop-down menu.
2. Click **Start Backup** in the upper right corner of the History page.
3. In the Start Backup dialog box, enter the following:

   Backup Name
   an optional name for the backup. The name must be unique. It is recorded in backup history and is displayed in the backup's Operation Details.

   Comment
   an optional free-form comment describing the backup. The comment is recorded in backup history and is displayed in the backup's Operation Details.

4. Click **OK** to start the backup.

   A notification is displayed when the backup starts and when it is completed.

TIP To see the status of the backup on the History page, refresh your browser.

TIP You can also use the sas-backup command to run a backup. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

Note: Recoveries cannot be run from SAS Backup Manager. Instead, use the sas-recover-offline command. Be sure to follow the instructions in "Performing a Recovery Using the Deployment Backup and Recovery Tool" in the *SAS Intelligence Platform: System Administration Guide*.

Viewing and Modifying the Backup Policy

View the Backup Policy

To view the current backup policy, select **Policy** from the drop-down menu. The Policy page displays the following:

Diagram (**Source View** and **Machine View**)
displays a tree diagram of the currently defined backup sources. To see a different view of the diagram:

- Click the Source View icon in the toolbar to display a node for each backup source. Under each backup source, a child node is displayed for each host machine for that source.
- Click the Machine View icon in the toolbar to display a node for each host machine. Under each machine, child nodes are displayed for the backup sources that are on the machine.

When a diagram is displayed, you can do the following:

- Zoom in or out by clicking the diagram to select it and then pressing **Ctrl** while scrolling the mouse wheel.
- If parts of the diagram are not visible, use the mouse to drag the entire diagram right, left, upward, or downward.
- Click a node to collapse its child nodes.
- Click the node again to expand it so that its child nodes reappear.

Configuration Details
displays details about the current backup configuration.

TIP You can also use the sas-display-backup-config command to display the backup policy. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

Update the Backup Policy

To update the backup policy:

1. On the **Policy** page, click the Edit icon in the toolbar.
2. While the Configuration Details pane is in Edit mode, you can specify the following:

   Schedule Allowed
   Select **YES** or **NO** to indicate whether scheduled backups are to run. If you select **NO**, the backup schedule is ignored.

   Vault Location (optional)
   If you have a homogeneous operating system environment, use this field to specify the path to a central, network-accessible location to store backups. If a

vault location is specified, backups are automatically copied from the various host machines to the vault following each backup operation. A vault is highly recommended to avoid the loss of backup files in the event that a host machine fails.

The path that you specify must meet these requirements:

- The path must already exist.
- All of the host machines that are included in the backup must be able to access the vault location using the specified path.
- You cannot specify a path that is currently used to store a host machine's local backups.
- On Windows machines, specify a UNC path.
- On Windows systems, make sure that the backup user has Read and Write access to the vault location. See "Specifying a Backup User to Access a Central Vault Location or a Clustered Metadata Server (Windows Only)" in the *SAS Intelligence Platform: System Administration Guide*.
- On UNIX, make sure the SAS Installer user for each server and middle-tier machine has Read and Write access to the central vault location. After granting permissions on UNIX, you must restart the SAS Deployment Agent.

Note: You can use this field to specify or change a vault location only if you have a homogeneous operating system environment. A homogeneous environment is one in which all of the host machines that are included in the backup are in the same operating system family. For example, Solaris and HP-UX machines are both considered to be in the UNIX operating system family.

Note: In heterogeneous (mixed) operating system environments, you must leave the **Vault Location** blank and use manual steps to specify the central vault. See "Specifying a Central Vault Location in a Heterogeneous Operating System Environment" in the *SAS Intelligence Platform: System Administration Guide*.

Days to retain backup (optional)
Enter the number of days that backup files are to be retained in the central vault location (if specified) and in the local backup location on each machine. Backups directories and their contents will be deleted automatically after this number of days. If no retention period is specified, backups are retained for 30 days by default.

Email List (optional)
Specify one or more email addresses to which alert messages are to be sent when a backup or a recovery is unsuccessful. After you enter an address, click the Add icon ➕ next to your entry.

Note: The following special characters are not supported in email addresses: ( ) : \ ; < > [}

To delete an address, click ⊗ next to the address.

If you do not specify any email addresses, alert messages are sent to the system administrator email address that was specified in the SAS Deployment Wizard.

3. When you are finished, click the Save icon 💾 in the toolbar.

TIP You can also use the sas-update-backup-config command to modify the backup policy. The command provides additional options including filtering physical data,

adding or removing tiers or database instances, and reorganizing metadata repositories. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

View Configuration Information for Backup Sources

Backup sources are discovered automatically. The sources are displayed in the Source View and Machine View diagrams, and they are also listed at the bottom of the Configuration Details pane. To view additional information about a source, click the Collapsed arrow ▶ to the left of the source name. The following information is displayed:

Host
the host name of the machine where the source is located.

Included
indicates whether the source is currently included or excluded from backups.

Note: This setting cannot be changed in the SAS Backup Manager user interface. To include or exclude a backup source, use the command sas-update-backup-config.

Operating System
the host name of the machine where the source is located.

Configurable Path
the path to the configuration directory for this source. This field is not applicable to all source types.

SAS Config
the path to the `Levn` directory that is associated with this backup source.

Includes and **Excludes**
lists any filters that are associated with this backup source.

The source information is for display only. To filter physical data or add or remove tiers, servers, or database instances from the backup configuration, use the sas-update-backup-config command. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

Viewing and Modifying the Backup Schedule

View the Backup Schedule

To view the backup schedule, select **Schedule** from the drop-down box. The Schedule page displays a row for each time of day that backups are scheduled to run. Check marks in the columns indicate the scheduled days of the week for each time.

Best Practices for Scheduling Backups

Follow these best practices when scheduling backups:

- If you schedule multiple backups per day, be sure to leave enough time for each backup job to complete before the next scheduled backup starts.

- To avoid conflicts, make sure not to schedule the deployment backup at the same time as the stand-alone SAS Metadata Server backups. By default, metadata server backups are scheduled to run at 1:00 a.m. server local time every day except Sunday.

  Note: To see the current metadata server backup schedule, go to the **Plug-ins** tab of SAS Management Console and select **Environment Management** ⇨ **Metadata Manager** ⇨ **Metadata Utilities** ⇨ **Server Backup** ⇨ **Backup Schedule**.

Enable or Disable Scheduled Backups

In order for schedules to be in effect, the **Schedule Allowed** backup option must be set to **YES**. This setting is turned on by default. In addition, make sure that the SAS Deployment Agent is running. See "Starting and Stopping the SAS Deployment Agent" in the *SAS Intelligence Platform: System Administration Guide*.

If you want to disable scheduled backups for some reason, change the **Schedule Allowed** setting to **NO**. See "Update the Backup Policy" on page 51.

Update the Backup Schedule

By default, the SAS Deployment Wizard schedules backups to be performed automatically each Sunday at 1:00 a.m. You can modify the schedule as follows:

- To change an existing row in the schedule, click the Edit icon in the toolbar. You can then do the following:
  - Specify different times by clicking the **Time** field. Use the time selector to specify the backup start time, and then click **OK**.
  - Specify different days of the week by selecting or clearing the check boxes in the columns **Mon** through **Sun**.
  - Delete the row by clicking the Delete icon in the first column.

  When you are finished, click **Save**.
- To add another row to the schedule:
  1. Click the Add icon ➕ in the toolbar.

     A new row is added to the schedule with the default time (1:00 a.m.) and default day (Sunday) selected.
  2. In the new row, click the **Time** field. Use the time selector to specify the backup start time, and then click **OK**.
  3. Select or clear the appropriate check boxes in the columns **Mon** through **Sun**.
  4. Click **Save**.

TIP You can also use the commands sas-set-backup-schedule, sas-list-backup-schedule, and sas-delete-backup-schedule to manage the backup schedule. See "Deployment Backup Command Reference" in the *SAS Intelligence Platform: System Administration Guide*.

Recommended Reading

Here is the recommended reading list for this title

- The online Help for SAS Environment Manager 2.5.
- *SAS Environment Manager: User's Guide*
- *SAS Intelligence Platform: System Administration Guide*
- *SAS Intelligence Platform: Security Administration Guide*
- *SAS Intelligence Platform: Data Administration Guide*
- *SAS Management Console: Guide to Users and Permissions*
- *SAS Guide to Metadata-Bound Libraries*

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

SAS Books
SAS Campus Drive
Cary, NC 27513-2414
Phone: 1-800-727-0025
Fax: 1-919-677-4444
Email: sasbook@sas.com
Web address: sas.com/store/books

Index

A

access control templates 35
ACTs
  See access control templates
administrators
  Administer permission 28
  intermittent 12
authentication domains 17

C

capabilities 13

D

default ACT
  See repository ACT

G

groups 13
  DBMS access 10

I

identity precedence 27
internal accounts 16

L

logins 14
  qualifying Windows user IDs in 14
  uniqueness requirement 19

M

members 14

P

passwords 18
permissions 28
  inheritance 27

R

repository ACT 26
roles 13

U

user administration 11
user IDs
  qualifying in logins 14
  uniqueness requirement 19
users 12
  DBMS access 10
  dual 12
  intermittent administrators 12

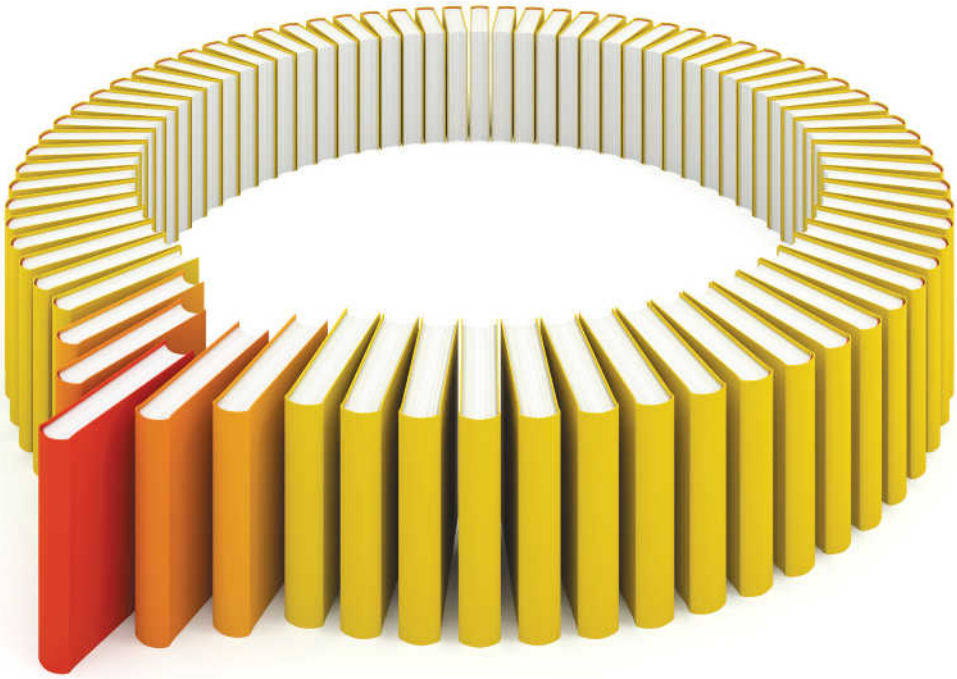

Gain Greater Insight into Your SAS® Software with SAS Books.

Discover all that you need on your journey to knowledge and empowerment.

support.sas.com/bookstore
for additional books and resources.

sas
THE POWER TO KNOW.

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 trademarks of their respective companies. © 2013 SAS Institute Inc. All rights reserved. S107969US.0613