

> PASW[®] Collaboration and Deployment Services 4.1 Administrator's Guide



SPSS Inc. 233 South Wacker Drive, 11th Floor
Chicago, IL 60606-6412
Tel: (312) 651-3000
Fax: (312) 651-3668

SPSS is a registered trademark.

PASW is a registered trademark of SPSS Inc.

The SOFTWARE and documentation are provided with RESTRICTED RIGHTS. Use, duplication, or disclosure by the Government is subject to restrictions as set forth in subdivision (c) (1) (ii) of The Rights in Technical Data and Computer Software clause at 52.227-7013. Contractor/manufacturer is SPSS Inc., 233 South Wacker Drive, 11th Floor, Chicago, IL 60606-6412.

Patent No. 7,023,453

Licensee understands and agrees that the Sample Code provided hereunder is provided as-is without warranty. Licensee further agrees that SPSS Inc. or its suppliers are not required to maintain or support such Sample Code. Licensee's right to use such code shall be set forth in a separate agreement between SPSS Inc. or a distributor of SPSS Inc. and Licensee.

General notice: Other product names mentioned herein are used for identification purposes only and may be trademarks of their respective companies.

Windows and Active Directory are registered trademarks of Microsoft Corporation in the United States and/or other countries.

OpenLDAP is a registered trademark of the OpenLDAP Foundation.

Eclipse is a registered trademark of the Eclipse Foundation. DataDirect, DataDirect Connect, INTERSOLV, and SequeLink are registered trademarks of DataDirect Technologies.

Copyright (c) 1995-2003 International Business Machines Corporation and others All rights reserved.

Permission is hereby granted, free of charge, to any person obtaining a copy of this software and associated documentation files (the "Software"), to deal in the Software without restriction, including without limitation the rights to use, copy, modify, merge, publish, distribute, and/or sell copies of the Software, and to permit persons to whom the Software is furnished to do so, provided that the above copyright notice(s) and this permission notice appear in all copies of the Software and that both the above copyright notice(s) and this permission notice appear in supporting documentation.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OF THIRD PARTY RIGHTS. IN NO EVENT SHALL THE COPYRIGHT HOLDER OR HOLDERS INCLUDED IN THIS NOTICE BE LIABLE FOR ANY CLAIM, OR ANY SPECIAL INDIRECT OR CONSEQUENTIAL DAMAGES, OR ANY DAMAGES WHATSOEVER RESULTING FROM LOSS OF USE, DATA OR PROFITS, WHETHER IN AN ACTION OF CONTRACT, NEGLIGENCE OR OTHER TORTIOUS ACTION, ARISING OUT OF OR IN CONNECTION WITH THE USE OR PERFORMANCE OF THIS SOFTWARE.

Except as contained in this notice, the name of a copyright holder shall not be used in advertising or otherwise to promote the sale, use or other dealings in this Software without prior written authorization of the copyright holder.

Printed in the United States of America.

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, recording, or otherwise, without the prior written permission of the publisher.

Preface

PASW Collaboration and Deployment Services enable widespread use and deployment of predictive analytics with features like centralized, secure, and auditable storage of analytical assets, advanced capabilities for management and control of predictive analytic processes, as well as sophisticated mechanisms of delivering the results of analytical processing to the end users.

This manual, , documents the administration aspects of operating PASW Collaboration and Deployment Services. The information is provided for tasks such as setting up content repository server, managing users, configuring communication protocols, installing updates, auditing repository, etc. Software and hardware requirements for PASW Collaboration and Deployment Services and the system installation and configuration are documented in *PASW Collaboration and Deployment Services 4.1 Installation and Configuration Guide*. The tasks associated with everyday use of the analytical facilities of PASW Collaboration and Deployment Services are documented in *Deployment Manager 4.1 User's Guide*.

Technical Support

The services of SPSS Inc. Technical Support are available to registered customers of SPSS Inc.. Customers may contact Technical Support for assistance in using SPSS Inc. products or for installation help for one of the supported hardware environments. To reach Technical Support, see the SPSS Inc. Web site at <http://www.spss.com>, or contact your local office, listed on the SPSS Inc. Web site at <http://www.spss.com/worldwide>. Be prepared to identify yourself, your organization, and the serial number of your system.

Tell Us Your Thoughts

Your comments are important. Please let us know about your experiences with SPSS Inc. products. Please send e-mail to suggest@us.ibm.com, or write to SPSS Inc., Attn: Director of Product Planning, 233 South Wacker Drive, 11th Floor, Chicago IL 60606-6412.

Contents

1	Overview	1
	PASW Collaboration and Deployment Services	1
	Collaboration	1
	Deployment	2
	System Architecture	2
	Repository	3
	Deployment Manager	4
	Deployment Portal	4
	Browser-based Deployment Manager	5
	Enterprise View	5
	Execution Servers	5
	PASW BIRT Report Designer	6
	Products with Collaboration	6
2	Getting Started	8
	Starting the Repository	8
	Stopping the Repository	11
	Using the Browser-based Deployment Manager	12
	Changing Passwords	13
	Navigating through the Browser-based Deployment Manager	13
	Accessing System Information	14
	Using the Deployment Manager	15
	Administered Servers	15
	Adding New Administered Servers	16
	Viewing Administered Server Properties	19
	Connecting to Administered Servers	19
	Disconnecting Administered Servers	20
	Deleting Administered Servers	20
	Naming Conventions	21
3	Users and Groups	22
	Setting Up PASW Collaboration and Deployment Services Users	22

Managing Users and Groups in Deployment Manager	23
Creating a User	26
Editing a User	28
Deleting a User	29
Creating a Group	29
Editing a Group	31
Deleting a Group	32
Importing Users and Groups	32
Creating an Extended Group	32
Creating an Allowed User	34
4 Roles	36
Roles Overview	36
Actions	36
Administrators Role	38
Manage Role Definitions	38
Creating a New Role	38
Editing a Role	39
Editing Users and Groups Assigned to a Role	40
Removing a Role	41
5 Security Providers	42
Security Providers in Deployment Manager	43
Configuring Security Providers	43
Security Providers in the Browser-based Deployment Manager	49
Configuring Security Providers	49
6 Single Sign-On	52
Configuring Single Sign-on	52
7 Repository Configuration	55
Administrator	56

Cache Provider	56
Custom Dialogs	57
Data Service	57
Deployment Manager	57
Deployment Portal	58
Deployment Portal Scoring	58
Enterprise View	59
Help	60
Notification	60
PASW BIRT Report Designer	64
Pager	64
Process Management	64
Reporting	66
Repository	66
Scoring Service	68
Search	69
Security	69
Setup	70
ShowCase	71
8 <i>MIME Types</i>	72
Adding MIME Type Mappings	73
Editing MIME Type Mappings	74
Deleting MIME Type Mappings	74
9 <i>Reindexing the Repository</i>	75
10 <i>Notifications</i>	77
Notification Message Template Structure	77
Message Properties	78
Message Content	80
Message Format	82
Editing Notification Templates	83
Job Status	84

Optimizing Notification Service Performance	87
Notification Service Configuration	88
General Recommendations	89
Debugging the Notification Service	91
Troubleshooting Notification Delivery Failures	91
11 JMS Setup	93
JMS Topic Configuration	93
PASW Collaboration and Deployment Services Setup	94
Sample JMS Client Program	94
Running the Sample Program	96
Message-Based Processing Example	96
12 Auditing the Repository	97
Database Audit Facilities	97
Audit Events	98
Event Tables	99
Audit Views	101
Audit (SPSSPLAT_V_AUDIT)	101
Custom Property (SPSSPLAT_V_CUSTOMPROPERTY)	102
File Version (SPSSPLAT_V_FILEVERSION)	102
Job History (SPSSPLAT_V_JOBHISTORY)	103
Job Step (SPSSPLAT_V_JOBSTEP)	104
Schedule (SPSSPLAT_V_SCHEDULE)	105
Stream Attribute Value (SPSSPLAT_V_STREAMATTRVALUE)	106
Stream Node (SPSSPLAT_V_STREAMNODE)	106
Scoring Service Logging	107
Request Log Table	107
Database Views	108
Audit Query Examples	110

Appendices

A Troubleshooting **113**

PASW Collaboration and Deployment Services	113
Solaris	115
HP-UX	116
Oracle Database.	116
JBoss.	117
WebLogic	117
WebSphere	118

B Nativestore Schema Reference **120**

nativestore Element	120
user Element	120
obsolete Element.	122

Index **124**

Overview

PASW Collaboration and Deployment Services

PASW Collaboration and Deployment Services is an enterprise-level application that enables widespread use and deployment of predictive analytics. PASW Collaboration and Deployment Services provides centralized, secure, and auditable storage of analytical assets and advanced capabilities for management and control of predictive analytic processes, as well as sophisticated mechanisms for delivering the results of analytical processing to the end users. The benefits of PASW Collaboration and Deployment Services include:

- safeguarding the value of analytical assets
- ensuring compliance with regulatory requirements
- improving the productivity of analysts
- minimizing the IT costs of managing analytics

PASW Collaboration and Deployment Services allows you to securely manage diverse analytical assets and fosters greater collaboration among those developing and using them. Furthermore, the deployment facilities ensure that the right people get the information they need to take timely, appropriate action.

Collaboration

Collaboration refers to the ability to share and reuse analytic assets efficiently, and is the key to developing and implementing analytics across an enterprise. Analysts need a location in which to place files that should be made available to other analysts or business users. That location needs a version control implementation for the files to manage the evolution of the analysis. Security is required to control access to and modification of the files. Finally, a backup and restore mechanism is needed to protect the business from losing these crucial assets.

To address these needs, PASW Collaboration and Deployment Services provides a repository for storing assets using a folder hierarchy similar to most file systems for organization. Files stored in the repository are available to users throughout the enterprise, provided those users have the appropriate permissions for access. To assist users in finding assets, the repository offers a search facility.

Analysts can work with files in the repository from client applications that leverage the service interface of PASW Collaboration and Deployment Services. Products such as PASW Statistics and PASW Modeler allow direct interaction with files in the repository. An analyst can store a version of a file in development, retrieve that version at a later time, and continue to modify it until it is finalized and ready to be moved into a production process. These files can include

custom interfaces that run analytical processes allowing business users to take advantage of an analyst's work.

The use of the repository protects the business by providing a central location for analytical assets that can be easily backed-up and restored. In addition, permissions at the user, file, and version label levels control access to individual assets. Version control and object version labels ensure the right versions of assets are being used in production processes. Finally, logging features provide the ability to track file and system modifications.

Deployment

To realize the full benefit of predictive analytics, the analytic assets need to provide input for business decisions. Deployment bridges the gap between analytics and action by delivering results to people and processes on a schedule or in real time.

In PASW Collaboration and Deployment Services, individual files stored in the repository can be included in processing **jobs** that define an execution sequence for the files. The execution results can be stored in the repository, on a file system, or delivered to specified recipients. Results stored in the repository can be accessed by any user with sufficient permissions using the Deployment Portal interface. The jobs themselves can be triggered according to a defined schedule or in response to system events.

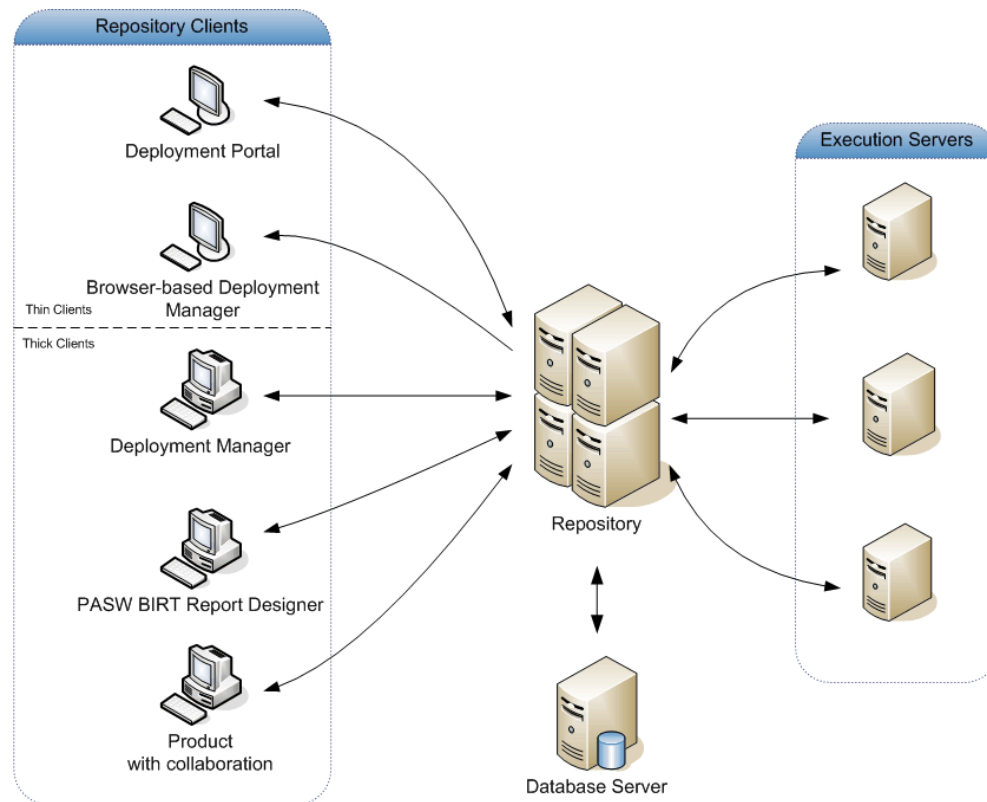
In addition, the scoring service of PASW Collaboration and Deployment Services allows analytical results from deployed models to be delivered in real time when interacting with a customer. An analytical model configured for scoring can combine data collected from a current customer interaction with historical data to produce a score that determines the course of the interaction. The service itself can be leveraged by any client application, allowing the creation of custom interfaces for defining the process.

The deployment facilities of PASW Collaboration and Deployment Services are designed to easily integrate with your enterprise infrastructure. Single sign-on reduces the need to manually provide credentials at various stages of the process. Moreover, the system can be configured to be compliant with Federal Information Processing Standard Publication 140-2.

System Architecture

In general, PASW Collaboration and Deployment Services consists of a single, centralized repository that serves a variety of clients, using execution servers to process analytical assets.

Figure 1-1
PASW Collaboration and Deployment Services Architecture



PASW Collaboration and Deployment Services consists of the following components:

- repository for analytical artifacts
- Product with Collaboration
- Deployment Manager
- Deployment Portal
- browser-based Deployment Manager
- Enterprise View
- PASW BIRT Report Designer

Repository

The repository provides a centralized location for storing analytical assets, such as models and data. The repository includes facilities for:

- Security
- Version control
- Searching
- Auditing

The repository requires an installation of a relational database, such as Oracle, IBM DB2 UDB, or Microsoft SQL Server.

Configuration options for the repository are defined using the Deployment Manager or the browser-based Deployment Manager. The contents of the repository are managed with the Deployment Manager and accessed with Deployment Portal.

Deployment Manager

Deployment Manager is a client application that allows users to schedule, automate, and execute analytical tasks, such as updating models or scores, using the repository. Deployment Manager allows a user to:

- View any existing files within the system, including reports, SAS syntax files, and data files.
- Import files into the repository.
- Schedule jobs to be executed repeatedly using a specified recurrence pattern, such as quarterly or hourly.
- Modify existing job properties in a user-friendly interface.
- Determine the status of a job.
- Specify e-mail notification of job status.

In addition, Deployment Manager allows users to perform administrative tasks for PASW Collaboration and Deployment Services, including:

- user management
- security provider configuration
- role and action assignment

Deployment Portal

Deployment Portal is a thin-client interface for accessing the repository. Unlike the browser-based Deployment Manager, which is intended for PASW Collaboration and Deployment Services administrators, Deployment Portal is a web portal serving a variety of users. Deployment Portal includes the following functionality:

- Browsing the repository content by folder
- Opening published content
- Running jobs and reports
- Generating scores using models stored in the repository
- Searching repository content.
- Viewing content properties.
- Accessing individual user preferences, such as e-mail address and password, general options, subscriptions, and options for output file formats.

Browser-based Deployment Manager

The browser-based Deployment Manager is a thin-client interface for performing setup and system management tasks, including:

- Configuring the system.
- Configuring security providers.
- Managing MIME types.

Non-administrative users can perform any of these tasks provided they have the appropriate actions associated with their login credentials. The actions are assigned by an administrator.

Enterprise View

Enterprise View provides a single, consistent view of enterprise data. Enterprise View allows users to define and maintain a common view of warehoused and transaction data needed to perform analytics, optimization, deployment, and reporting. Underlying data may come from a variety of sources, including a data warehouse, an operational data store, and an online transaction database. Enterprise View ensures a consistent use of enterprise data and hides the complexities of stored data structures from the end user. Enterprise View is the data backbone for the predictive enterprise.

Data discovery requires a major investment of resources from the organizations deploying predictive analytics. The process is labor intensive—it can involve representatives from departments across the organization and often entails resolving differences in data structure and semantics across organizational boundaries. Enterprise View provides a mechanism for recording the outcomes of the data discovery process, versioning and securing the resulting schema, and tracking changes over time.

Enterprise View includes the Enterprise View Driver component designed to provide other applications access to Enterprise View objects stored in the repository. The driver operates similarly to ODBC drivers with the exception that it does not directly query a physical data source but rather references Enterprise View data provider definitions and application views. Note that while Enterprise View is installed as part of Deployment Manager, Enterprise View Driver must be installed separately. For more information, see the installation instructions.

Execution Servers

Execution servers provide the ability to execute resources stored within the repository. When a resource is included in a job for execution, the job step definition includes the specification of the execution server used for processing the step. The execution server type depends on the resource.

Execution servers currently supported by PASW Collaboration and Deployment Services include:

- **SAS.** The SAS execution server is the SAS executable file *sas.exe*, included with Base SAS® Software. Use this execution server to process SAS syntax files.
- **Remote Process.** A remote process execution server allows processes to be initiated and monitored on remote servers. When the process completes, it returns a success or failure message. Any machine acting as a remote process server must have the necessary infrastructure installed for communicating with the repository.

Execution servers that process other specific types of resources can be added to the system by installing the appropriate adapters. For information, consult the documentation for those resource types.

During job creation, assign an execution server to each step included in the job. When the job executes, the repository uses the specified execution servers to perform the corresponding analyses.

PASW BIRT Report Designer

The reporting functionality of PASW Collaboration and Deployment Services is enabled by BIRT (Business Intelligence and Reporting Tools), an open-source package distributed by Eclipse Foundation under the Eclipse Public License. BIRT provides core reporting features, such as report layout, data access, and scripting. For more information about BIRT, see the [BIRT project page \(http://www.eclipse.org/birt\)](http://www.eclipse.org/birt).

The PASW Collaboration and Deployment Services installation includes the BIRT reporting engine server components, which enable the execution of BIRT report syntax files as part of the PASW Collaboration and Deployment Services reporting job steps. PASW BIRT Report Designer is a standalone application that can be used in conjunction with PASW Collaboration and Deployment Services. It provides a rich user interface with a number of advanced features for creating reports and must be installed separately.

If a PASW BIRT Report Designer report requires a JDBC-based database connection, a corresponding JDBC driver must be installed with the repository. For application server-specific information on the location of the JDBC drivers, see the corresponding section of the repository installation instructions.

To start PASW BIRT Report Designer, execute the file *BIRT.exe* in the installation directory. For information on using PASW BIRT Report Designer, see the documentation installed with the application.

Products with Collaboration

A product with collaboration allows interaction with the repository from within the native interface. Files can be stored and retrieved directly from the collaborating product.

In addition, some files stored in the repository can be executed as steps within jobs. A job can contain any number of steps, with each step corresponding to a separate file. Relationships defined between the steps determine the processing flow. The job can be scheduled to execute at a specific time, according to a recurrence pattern, or in response to a defined event. Moreover, notifications can be sent to specified recipients to report on individual step and overall job execution status.

Collaboration between PASW Collaboration and Deployment Services and other products is enabled through the use of adapters. These adapters are installed into the PASW Collaboration and Deployment Services environment to add the product-specific features. For more information, consult the collaborating product documentation.

Getting Started

After successfully installing repository, the following actions can be performed:

- Starting the server as a console application or service
- Stopping the server as a console application or service
- Logging in to and out of the system
- Changing passwords and navigating the interface
- Adding or changing PASW Modeler support

Starting the Repository

The repository can run either on a console or in the background. Running on a console allows viewing of processing messages and can be useful for diagnosing unexpected behavior. However, typically the repository runs in the background, handling requests from clients, such as PASW Modeler or the Deployment Manager.

Note: Running other applications simultaneously may reduce system performance and startup speed.

On the Windows platform, running on a console corresponds to running in a command window. Running in the background corresponds to running as a Windows service. In contrast, on a UNIX platform, running on a console corresponds to running in a shell, and running in the background corresponds to running as a daemon.

Note: To avoid permissions conflicts on UNIX systems, repository must always be started under the same credentials, preferably *root*.

The repository can be started using the scripts provided with the installation or native application server administration tools, such as WebLogic Server Administration Console. For more information, see the application server vendor documentation.

Starting as a Service under Windows

For an installation using JBoss, to start the repository as a service under Windows:

1. Open the Windows Services console.
2. Select PASW Platform Server from the list of available services.
3. Set the Startup Type to Automatic. This allows the repository to start when the system boots.
4. If not currently active, click Start.

Note: It may take several minutes for the server and its components to become active.

5. Click OK to accept the changes and close the dialog box.

For an installation using Weblogic, the *SPSSDomain* server needs to be set up as a Windows service. For additional information, see [WebLogic documentation \(http://e-docs.bea.com/wls/docs81/adminguide/winservice.html\)](http://e-docs.bea.com/wls/docs81/adminguide/winservice.html).

For WebSphere autostart configuration information, see [WebSphere documentation \(http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp\)](http://publib.boulder.ibm.com/infocenter/wasinfo/v6r1/index.jsp).

When you are using the OracleAS installation, use OPMN (Oracle Process Management and Notification) service to set up PASW Collaboration and Deployment Services autostart. OPMN itself is installed as a Windows service by default, so it can be started automatically by Windows. For more information, see [OracleAS documentation \(http://download.oracle.com/docs/cd/B31017_01/core.1013/b28944/toc.htm\)](http://download.oracle.com/docs/cd/B31017_01/core.1013/b28944/toc.htm).

Starting as a Daemon under UNIX

For information on starting the repository daemon on an application server other than JBoss, see the application server vendor documentation.

To start the repository as a daemon under UNIX for an installation using JBoss:

1. From a command line, navigate to the *bin* directory of the JBoss installation.
2. Enter:

```
./mm.sh start
```

Note: It may take several minutes for the server and its components to become active.

3. To verify that the server is active, enter `ps -e | grep wrapper`. If active, the process appears in the list.

Starting on IBM i

The WebSphere instance running PASW Collaboration and Deployment Services on IBM i is started by executing commands similar to the following in QShell environment:

```
cd /QIBM/UserData/WebSphere/AppServer/V61/Base/profiles/<profile name>/bin
./startServer <profile name>
```

To configure application servers to start automatically when the QWAS61 (WebSphere Application Server subsystem) starts:

1. Give your user profile authority to the QWAS61/QWASJOB job description and QWAS61/QWAS61 subsystem description.

2. For each profile, create a duplicate of the job description used by WebSphere Application Server profiles. Use the following command on the command line:

```
CRTDUPOBJ OBJ(QWASJOB) FROMLIB(QWAS61) OBJTYPE(*JOB) TOLIB(mywasjobd) NEWOBJ(myserver)
```

3. Use the `CHGJOB` command to change the newly created job description such that the Request data or command (`RQSDTA`) field starts the new server. For example, to start the default profile's application server (`server1`) when the subsystem is started, set the `RQSDTA` field as follows:

```
'QSYS/CALL PGM(product_library/QWASSTRSVR) PARM("-profilePath" "user_data_root/profiles/default"
"-server""server1")'
```

4. Add an autostart job entry to the QWAS61/QWAS61 subsystem. Enter the following command from the CL command line:

```
ADDAJE SBS(QWAS61/QWAS61) JOB(myserver) JOB(mywasjobd/myserver)
```

Optional: Configure the system such that the QWAS61 subsystem starts during system startup. To enable automatic startup, add the following line to the system startup program: `STRSBS QWAS61/QWAS61`.

Note: The system startup program is defined by the `QSTRUPPGM` system value. TCP/IP must be active before WebSphere Application Server subsystem can start. Ensure that the `STRTCP` command runs before the `STRSBS QWAS61/QWAS61` command in your startup program or in your autostart job.

Note: These commands can be added to any startup and shutdown rc scripts.

Starting as a Console Application under Windows

To start repository as a console application under Windows:

1. If JBoss is being used, from the repository installation path, execute `startserver.bat`. Alternatively, if Weblogic is being used, from the `<BEA HOME>/user_projects/domains/<PASW Collaboration and Deployment Services domain>/` subdirectory of the repository installation path, execute `startWebLogic.cmd`. For both application servers, the batch file opens a command line window and starts the server.

Note: It may take several minutes for the server and its components to fully load.

2. Verify that the server is active by opening a browser. The login screen appears when the IP address and port number are entered in the URL field.

Starting as a Console Application under UNIX

To start the repository as a console application under UNIX:

1. If JBoss is being used, from the command line, navigate to the repository installation directory and execute `startserver.sh`. Alternatively, if Weblogic is being used, navigate to the `<BEA HOME>/user_projects/domains/<PASW Collaboration and Deployment Services domain>/` subdirectory of the repository installation and execute `startWebLogic.sh`. In both cases, the shell script runs and starts the server.

Note: It may take several minutes for the server and its components to fully load.

2. To verify that the server is active, enter `ps -e | grep wrapper`. If active, the process appears in the list.

Stopping the Repository

The repository is halted from the command line console or stopped from the service menu.

Stopping the Service under Windows

To halt an active instance of the repository:

1. Open the Windows Services console.
2. Select PASW Platform Server from the list of available services.
3. Click Stop to halt the service.

Note: It may take several minutes for the server and its components to completely stop.

Stopping the Daemon under UNIX

For information on stopping the repository daemon on an application server other than JBoss, see the application server vendor documentation.

To halt an active instance of the repository for an installation using JBoss:

1. From a command line, navigate to the *bin* directory of the JBoss installation.
2. Enter:

```
./mm.sh stop
```
3. To verify that the server has stopped, enter `ps -e | grep wrapper`. The process should not appear in the list.

Stopping on IBM i

The WebSphere instance running PASW Collaboration and Deployment Services on IBM i is stopped by executing commands similar to the following in QShell environment:

```
cd /QIBM/UserData/WebSphere/AppServer/V61/Base/profiles/<profile name>/bin
./stopServer <profile name>
```

Stopping the Console Application under Windows

To halt an active instance of the repository:

1. Open the corresponding command line window and issue the Ctrl-C keystroke combination. This signals the command line to terminate the process.
2. Enter Y when prompted.

Stopping the Console Application under UNIX

To halt an active instance of the repository:

1. From a command line, enter `ps -e | grep wrapper`.
2. Use the `kill` command to terminate the corresponding PID.

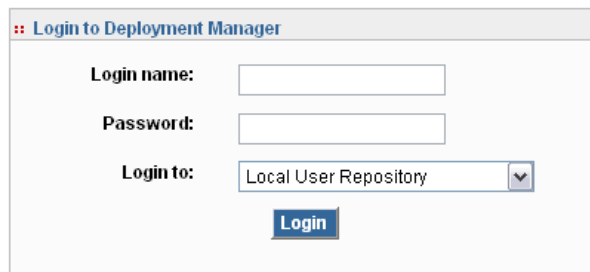
Note: It may take several minutes for the server and its components to completely stop.

Using the Browser-based Deployment Manager

The Login page is your gateway into the system. To log in:

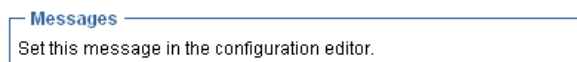
- ▶ Navigate to the Login page. Typically, the URL is `http://<host IP address>:<port number>/security/login`. The Login page opens. Note that use of *localhost* in place of the IP address may fail for some application servers; use of the IP address is recommended in all cases.

Figure 2-1
Login dialog box



The screenshot shows a web browser window titled "Login to Deployment Manager". Inside the window, there are three input fields: "Login name:" with an empty text box, "Password:" with an empty text box, and "Login to:" with a dropdown menu currently showing "Local User Repository". Below these fields is a blue "Login" button.

[Change Password?](#)



The screenshot shows a "Messages" section with a blue border. Inside, there is a single message: "Set this message in the configuration editor."

- ▶ In the Login Name field, enter your user ID.
- ▶ In the Password field, enter your password.
- ▶ Click Login. By default, the Configuration page appears.

Additional Options

On the Login page, you also have the option of changing your password. For more information, see the topic [Changing Passwords](#) on p. 13.

Important! Browser-based Deployment Manager does not allow single-sign on.

Changing Passwords

To change your password:

On the Login page, click Change Password? The Change Password dialog box opens.

Figure 2-2

Changing your password

The screenshot shows a web-based dialog box titled "Change Password". It contains the following elements:

- Login name:** A text input field containing the value "admin".
- Current Password:** A password input field containing four asterisks "****".
- New Password:** A password input field containing seven asterisks "*****".
- Confirm New Password:** A password input field containing seven asterisks "*****".
- Save New Password:** A blue button located below the password fields.
- Return to Login:** A blue hyperlink located at the bottom right of the dialog.
- Messages:** A section at the bottom of the dialog containing the text: "Having problems with login? Contact system administrator: [Administrator](#) Change in [config editor](#)".

- ▶ In the Login Name field, enter your login name.
- ▶ In the Current Password field, enter your current password.
- ▶ In the New Password field, enter your new password.
- ▶ In the Confirm New Password field, reenter your new password.
- ▶ Click Save New Password. In the Messages section, the following text appears:
Password updated
- ▶ Click Return to Login. The Login page opens. Log in to the system using your new password. For more information, see the topic [Using the Browser-based Deployment Manager](#) on p. 12.

Navigating through the Browser-based Deployment Manager

The browser-based Deployment Manager relies primarily on tab-based navigation. In general, components of the system are organized from the general to the specific. From the navigation panel on the left, you can choose any of the following categories:

- Configuration
- Deployment Portal
- MIME Types
- Repository Index
- Security Providers

- Logout
- About
- Administrator Guide
- Help

Each of these items has one or more sections associated with it. When you click on an item, its corresponding section appears in the right-hand pane. If a section has multiple subsections, a series of tabs appears in the right-hand pane. By default, the contents of the first tab are displayed. For example, if you click MIME Types from the navigation list, the MIME Types and File Type Icons section appears.

Clicking Set versus Pressing Enter

The system is mouse-driven. It is not recommended that you use the Enter key to complete actions. Typically, pressing Enter will not submit your request. For example, throughout the system you will see the Set key. If you press Enter instead of clicking Set, your request will not be processed. Clicking Set commits your changes to the database.

Accessing System Information

The information about your PASW Collaboration and Deployment Services installation can be accessed using the About page. The page displays the version number for the system and also lists the information for individual components (installed packages), including the general component category (“Area”), version number, and license. The page also allows you to display detailed information listing the files included in each package, and provides the ability to download system information, installation logs, and application server logs. Application server logs can be used in troubleshooting the system. For more information, see the topic [Troubleshooting](#) in Appendix A on p. 113.

To display detailed information for installed packages

- ▶ Click Show Details.

To download a text file version of system information

- ▶ Click Download version and system details at the bottom of the page.

To download a zipped archive containing text files of version and system information and application server log

- ▶ Click Download version, system details and logs in one zip file at the bottom of the page. The files are downloaded as a zipped archive.

Using the Deployment Manager

Administrative tasks can be performed using the Deployment Manager as well as the browser-based Deployment Manager. An administrator can:

- Configure and enable security providers.
- Create users and groups for accessing the system.
- Define roles to control access to system features.

In addition, Deployment Manager allows administration of other servers, such as PASW Statistics and PASW Modeler servers.

Administered Servers

Server administration in Deployment Manager involves:

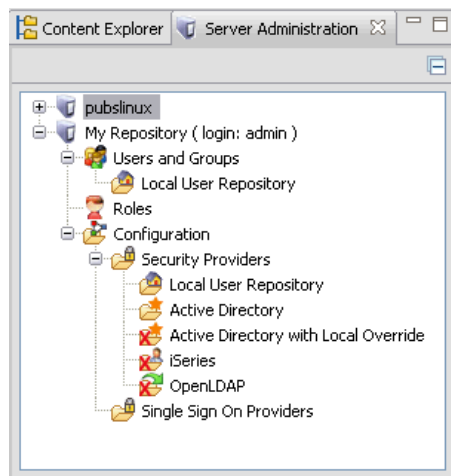
1. Adding the server to be administered to the system.
2. Logging in to the server being administered.
3. Performing administrative tasks for the server as needed.
4. Logging off from the server being administered.

The Server Administration tab offers access to this functionality. This tab lists the servers currently available to be administered. This list persists across Deployment Manager sessions, facilitating access to those servers.

From the menus choose:

Tools
Server Administration

Figure 2-3
Administered server list



The administered server list may include a variety of server types, including repository servers, PASW Modeler servers, and PASW Statistics servers. The actual administrative functionality available for a server depends on the server type. For example, security providers can be configured and enabled for repository servers but not for PASW Modeler servers.

Adding New Administered Servers

Before performing administrative tasks, a connection to the administered server must be established.

From the menus choose:

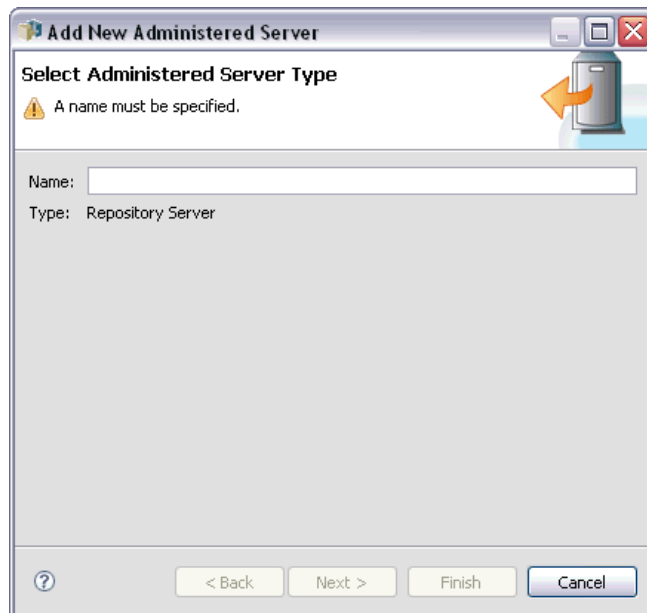
File
New
Administered Server Connection

The Add New Administered Server dialog box opens. Adding a new connection requires the specification of the administered server type and the administered security server information.

Selecting the Administered Server Name and Type

The first step of adding a new administered server to the system involves the definition of two parameters for the server—the name and the type.

Figure 2-4
Select Administered Server Type dialog box



Name. A label used to identify the server on the Server Administration tab. Including the port number in the name, such as *my_server:8080*, may help to identify the server in the administered server list.

Note: Alphanumeric characters are recommended. The following symbols are prohibited:

- quotation marks (single and double)
- ampersands (&)
- less-than (<) and greater-than (>) symbols
- periods
- commas
- semicolons

Type. The type of server being added. The list of possible server types depends on the system configuration and may include:

- Repository Server
- Administered PASW Modeler Server
- Administered PASW Statistics Server
- Administered PASW Text Analytics Server

Selecting an Administered Server Type


In the Select Administered Server Type dialog box:

1. Enter a name for the server.
2. Select the server type.
3. Click Next. The Administered Security Server Information dialog box opens.

Administered Server Information

The second step of adding a new administered server to the system involves the definition of the server properties.

Figure 2-5
Administered Security Server Information dialog box



Host. The name or IP address of the server.

Note: Alphanumeric characters are recommended. The following symbols are prohibited:

- quotation marks (single and double)
- ampersands (&)
- less-than (<) and greater-than (>) symbols
- periods
- commas
- semicolons

Port. The port number used for the server connection.

This is a secure port. Enables or disables the use of a Secure Sockets Layer (SSL) for the server connection. This option is not offered for all types of administered servers.

Specifying Administered Server Information

In the Administered Security Server Information dialog box:

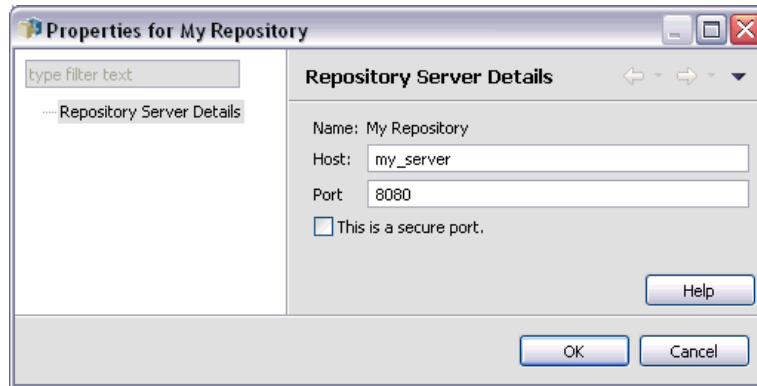
1. Enter the host name or IP address for the server being added.
2. Enter the port number on which the server being added is running.
3. Specify whether or not the server uses SSL, if applicable.
4. Click Finish.

The server appears in the administered server list on the Server Administration tab.

Viewing Administered Server Properties

To view the properties of an existing administered server, right-click the server on the Server Administration tab and select Properties from the drop-down menu. The Properties dialog box opens. The displayed properties depend on the type of server selected.

Figure 2-6
Repository Server properties



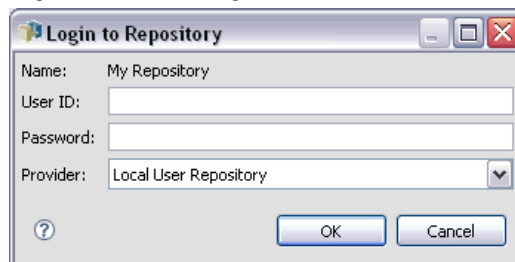
For repository servers, the properties include:

- **Label.** The name associated with the server as it appears on the Server Administration tab.
- **Host.** The name or IP address of the server.
- **Port.** The port number used for the server connection.
- **This is a secure port.** If this option is selected, the server uses an SSL connection for communication.

Connecting to Administered Servers

For most servers, you must connect to a server in the administered server list to perform administrative tasks. From the Server Administration tab, double-click the server to administer. The Login to Server dialog box opens.

Figure 2-7
Login to Server dialog box



For repository servers, the login parameters include:

- **User ID.** The user to log in to the server, displayed in clear text.

Password. The string used to authenticate the user. For security, password text is displayed in a masked format.

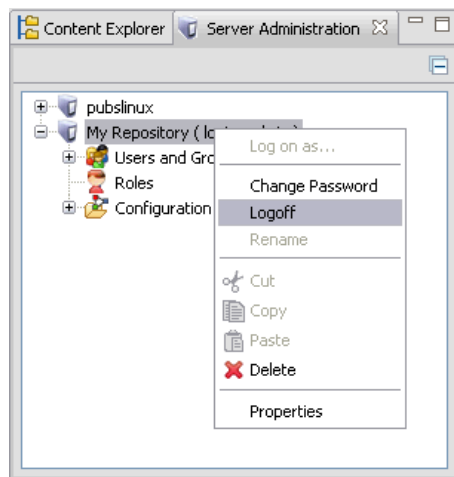
Provider. The provider against which to validate the specified login/password combination. This field appears only if multiple security providers are enabled for the system. Otherwise, the system validates the supplied credentials against the local user repository.

Disconnecting Administered Servers

After completing the desired administrative tasks, log off from the server.

1. On the Server Administration tab, right-click the server.
2. Select Logoff.

Figure 2-8
Logging off from a server



To administer the server, you must log in again.

Deleting Administered Servers

A server appears in the administered server list until it is deleted from the list.

1. On the Server Administration tab, select the server to delete.
2. From the menus choose:
Edit
Delete

Alternatively, right-click the server and select Delete from the drop-down menu.

If further administrative tasks for the server are needed in the future, the server will need to be added to the system again.

Naming Conventions

Throughout the system, you are asked to name entities, ranging from folders to topics. For example, you might want to add a new user or create a new topic.

The following naming conventions apply:

- Most characters, including spaces, are accepted by the system. However, the forward slash (/) is not allowed. If you type the forward slash as part of a name, it is not included in the name.
- The maximum character length is 255, including spaces.
- Names are not case sensitive.

Users and Groups

A PASW Collaboration and Deployment Services user is an individual or a process that is allowed to access files and execute programs. The user is authenticated with a user name and password pair against an internal or external database. Users have different levels of access to application resources.

Users can be organized into groups based on the need for information access and manipulation. Organizing users into groups helps minimize the effort required to distribute permissions to multiple users in a uniform and organized way.

Users and groups are assigned access to system resources through the mechanism of *roles*. A role is a set of actions predefined within the system, such as access to files and MIME types, ability to change system configuration, etc. Role assignments can be added or removed, and new roles can be established as needs change. Note that roles must be explicitly assigned before users can access the system. For more information, see the topic [Roles Overview](#) in Chapter 4 on p. 36.

PASW Collaboration and Deployment Services users and groups are handled by *security providers*. A security provider is the system that authenticates user credentials. Users and groups can be defined locally (in which case, PASW Collaboration and Deployment Services itself is the security provider) or derived from a remote directory, such as Windows Active Directory or OpenLDAP. For more information, see the topic [Security Providers](#) in Chapter 5 on p. 42.

Some environments may require setting up groups of remotely defined users that are specific to Deployment Manager. This will be the case if the groups specified in the remote directory are not fine-grained enough. The directory administrator may not be able to add these more specific groups because of policy restrictions or because queries of the remote directory from external applications may not be allowed. In these instances, locally specified groups of remote users, referred to as *extended groups*, will be added to the list of groups already defined in the remote directory.

In many environments, the number of users that exists in a remote directory is quite large, while only a small subset of the total user pool actually needs access to PASW Collaboration and Deployment Services. In this case, the administrator can specify a list of *allowed users*, and only those users will be allowed to log in. The allowed list acts as a filter on the user name, but the actual authentication of the user is performed against the remote directory in a normal manner.

Setting Up PASW Collaboration and Deployment Services Users

Local user setup in PASW Collaboration and Deployment Services involves:

1. Creating the user and, if necessary, assigning her to groups. Local user and groups can be managed through either Deployment Manager.
2. Defining the user's level of access by assigning the role on an individual or group basis. For more information, see the topic [Editing Users and Groups Assigned to a Role](#) in Chapter 4 on p.

40. If the role with an appropriate set actions does not exist, it must be established. For more information, see the topic [Creating a New Role](#) in Chapter 4 on p. 38.

Externally defined user setup in PASW Collaboration and Deployment Services involves:


1. Setting up the external security provider, if it has not yet been defined. The user will be derived from that security provider. For more information, see the topic [Configuring Security Providers](#) in Chapter 5 on p. 43.
2. Creating allowed users if access must be limited to a subset of the Active Directory with Local Override users. Allowed users can be created only with Deployment Manager.
3. Defining the extended group and adding the user to the group if the Active Directory with Local Override user must be assigned to a group that does not exist in the remote directory. Extended groups can be created only with Deployment Manager.
4. Assigning the role on an individual or group basis. Roles are assigned to remotely defined users in the same manner in which they are assigned to local users.

Managing Users and Groups in Deployment Manager

Deployment Manager allows you to manage local users and groups and allowed user and extended groups defined for the Active Directory with Local Override security provider. Prior to performing any actions with users or groups, navigate to the administrative interface that controls these areas.

1. From the Tools menu, choose Server Administration.
2. On the Server Administration tab, log in to a repository server. Double-click the Users and Groups icon to expand the hierarchy. If no external security providers are set up, Local User Repository is the only entry in the hierarchy. If Active Directory with Local Override has been configured as a security provider with allowed users or extended groups options enabled, the Active Directory with Local Override entry is also displayed.

Figure 3-1
Server Administration tab




3. Double-click the Local User Repository icon or Active Directory with Local Override icon.

The Manage Users and Groups editor opens.


- For Local User Repository, the editor displays all native users and groups or shows a filtered list based on the initial letters in the user and group names. An administrator can create and delete users and groups, edit the properties of existing users and groups, and import users and groups.

Figure 3-2
Manage Groups and Users editor



- For Active Directory with Local Override, the editor displays all externally defined groups and users that have been set up to access PASW Collaboration and Deployment Services or shows a filtered list based on the initial letters in the user and group names. An administrator can create and delete allowed users and extended groups and edit the properties of existing groups if the allowed users and extended groups options are enabled for the security provider. For more information, see the topic [Security Providers](#) in Chapter 5 on p. 42.


Figure 3-3
Manage Groups and Users editor for Active Directory with Local Override



Creating a User

In the Manage Users and Groups editor for Local User Repository, click New User. The Create New User dialog box opens.

Figure 3-4
Create New User dialog box



User Name. The name is not case-sensitive and can contain spaces.

Password. The local user's password. The password is case-sensitive.

Verify. Password verification field. If the passwords do not match, a message is displayed.

Show all available groups. Returns a list of all groups recognized by the system. Note that for very large directories there may be a limit on the number of entries that can be displayed. Therefore, it is recommended to specify a search string.

Show groups starting with. Filters the list of available groups according to the string entered. Use this field to refine the list of available groups.

Available groups. Lists the recognized groups to which the user can be assigned.

User belongs to groups. Lists the groups to which the user is currently assigned.

Add all. Associates all groups with the user.

Remove all. Disassociates all displayed groups from the user.

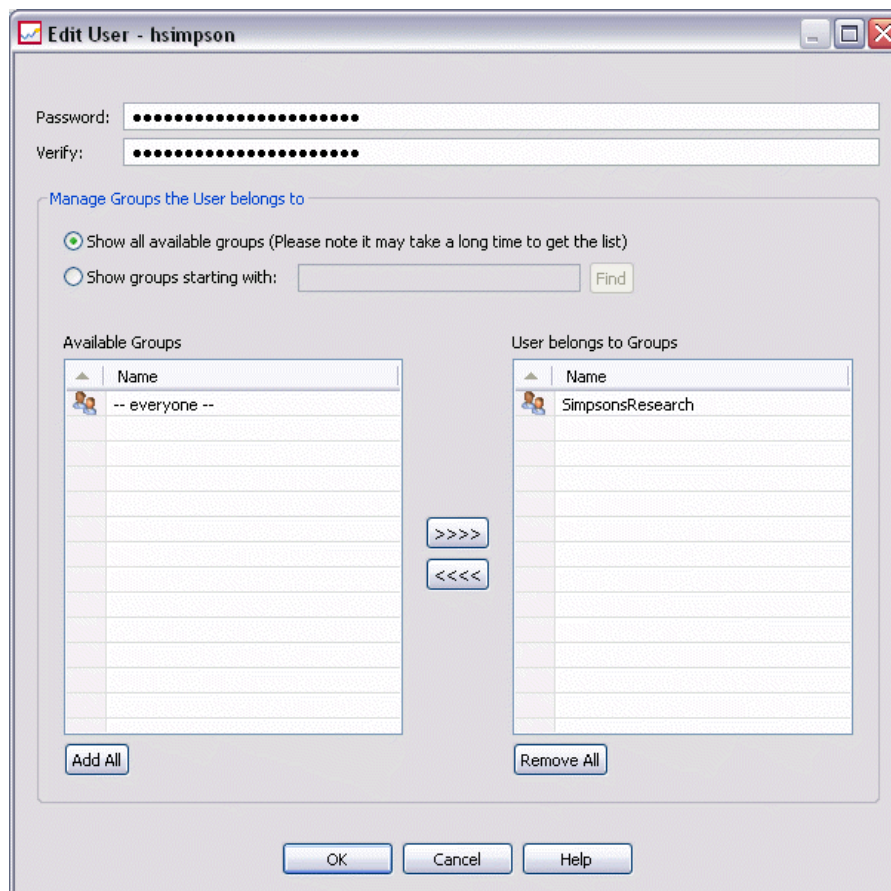
Creating a local user requires the login credentials to be specified. The user can also be associated with groups.

1. In the Create New User dialog box, specify the user name.
2. Specify the password.
3. Verify the password
4. If necessary, associate the user with groups.
5. Click OK. The new user appears in the list in the Manage Users and Groups editor.

Editing a User

Group assignments can be edited for local users and allowed users in Active Directory with Local Override. For local users, the password can also be edited. In the Manage Users and Groups editor, select the user and click Edit. The Edit User dialog box opens.

Figure 3-5
Edit User dialog box



Password. The local user's password. The password is case-sensitive.

Verify. Password verification field. If the passwords do not match, a message is displayed.

Show all available groups. Returns a list of all groups recognized by the system. Note that for very large directories there may be a limit on the number of entries that can be displayed. Therefore, it is recommended to specify a search string.

Show groups starting with. Filters the list of available groups according to the string entered. Use this field to refine the list of available groups.

Available groups. Lists the recognized groups to which the user can be assigned.

User belongs to groups. Lists the groups to which the user is currently assigned.

Add all. Associates all groups with the user.

Remove all. Disassociates all displayed groups from the user.

Deleting a User

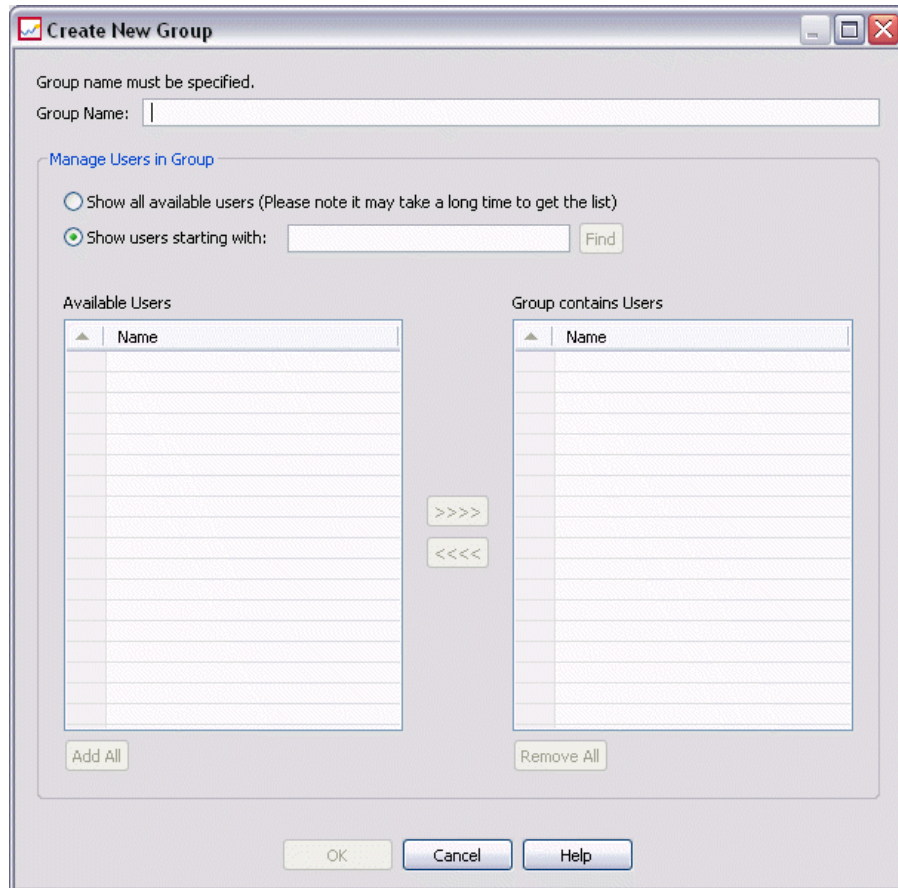
To delete a local user or an allowed user in Active Directory with Local Override:

1. Select the user in the Manage Users and Groups editor.
2. Click the Delete button. A dialog box opens to confirm that the user should be deleted.
3. Click Yes to delete the user from the system. The user is removed from the User/Group listing.

Creating a Group

In the Manage Users and Groups editor for Local User Repository, click New Group. The Create New Group dialog box opens.

Figure 3-6
Create New Group dialog box



Group Name. The name is not case-sensitive and can contain spaces.

Show all available users. Returns a list of all users recognized by the system. Note that for very large directories there may be a limit on the number of entries that can be displayed. Therefore, it is recommended to specify a search string.

Show users starting with. Filters the list of available groups according to the string entered. Use this field to refine the list of available groups.

Available users. Lists the recognized users that can be added to the group.

Group contains users. Lists the users assigned to the group.

Add all. Associates all users with the group.

Remove all. Disassociates all displayed users from the group.

Creating a local group requires the user name to be specified. Users can also be added to the group.

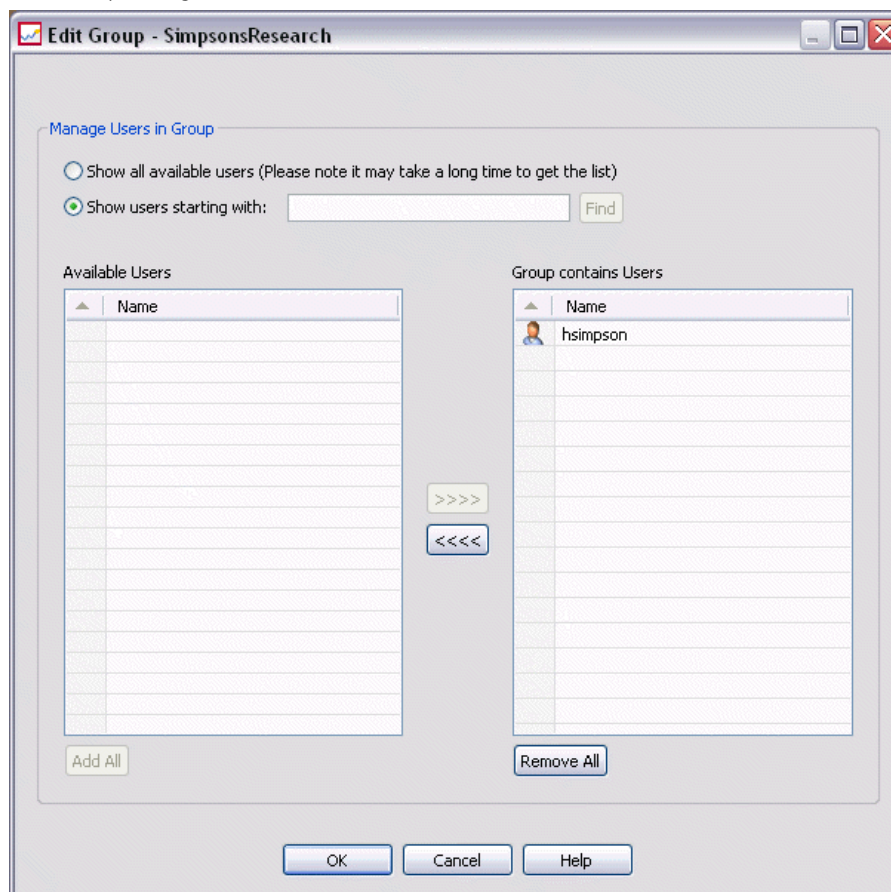
1. Specify the group name.
2. If necessary, add users to the group.

3. Click OK. The new group appears in the list in the Manage Users and Groups editor.

Editing a Group

The user list can be changed for local groups and extended groups in Active Directory with Local Override. In the Manage Users and Groups editor, select a group and click Edit. The Edit Group dialog box opens.

Figure 3-7
Edit Group dialog box



Show all available users. Returns a list of all users recognized by the system. Note that for very large directories there may be a limit on the number of entries that can be displayed. Therefore, it is recommended to specify a search string.

Show users starting with. Filters the list of available groups according to the string entered. Use this field to refine the list of available groups.

Available users. Lists the recognized users that can be added to the group.

Group contains users. Lists the users assigned to the group.

Add all. Associates all users with the group.

Remove all. Disassociates all displayed users from the group.

Deleting a Group

To delete a local group or an extended group in Active Directory with Local Override:

1. Select the group to delete in the Manage Users and Groups editor.
2. Click the Delete button. A dialog box opens to confirm that the entry should be deleted.
3. Click Yes to delete it from the system. The group is removed from the User/Group listing.


Importing Users and Groups

If you have to define a large number of local users or groups, you can use a principals import file to import users and groups in bulk. This file must follow the structure defined in the *nativestore.xsd* schema. For more information, see [Appendix B](#).

To import users and groups:

1. Click the Import button in the Manage Users and Groups editor for Local User Repository. The Import Users and Groups from File dialog box opens.

Figure 3-8
Import Users and Groups from File dialog box

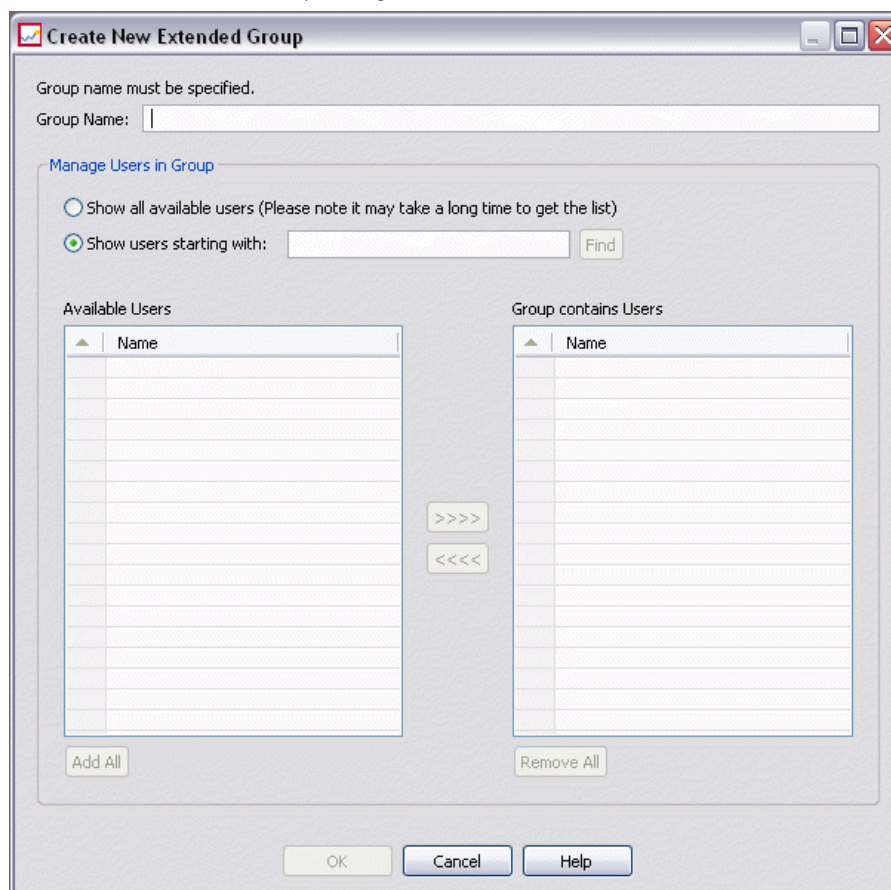


2. Select Update Users and Groups or Replace All Users and Groups.
 - **Update Users and Groups.** Updates the existing users with the information in the import file. Existing users and groups that are not defined in the file are not updated.
 - **Replace Users and Groups.** Replaces current users and groups with the information from the import file. Existing users and groups that are not defined in the file are removed.
3. Navigate to the location of the import file.
4. Click OK to import the file. The new users and groups appear in the list in the Manage Users and Groups editor.

Creating an Extended Group

In the Manage Users and Groups editor for Active Directory with Local Override, click New Extended Group. The Create New Extended Group dialog box opens.

Figure 3-9
Create New Extended Group dialog box



Show all available users. If allowed users option is enabled, returns the list of all allowed users. If allowed users option is disabled, a list of all users in the directory is returned. Note that for very large directories there may be a limit on the number of entries that can be displayed. Therefore, it is recommended to specify a search string.

Show users starting with. Filters the list of available groups according to the string entered. Use this field to refine the list of available groups.

Available users. Lists the recognized users that can be added to the group.

Group contains users. Lists the users assigned to the group.

Add all. Associates all users with the group.

Remove all. Disassociates all displayed users from the group.

Creating an extended group requires the user name to be specified. Users can also be added to the group.

1. Specify the group name.
2. If necessary, add users to the group.

3. Click OK. The new extended group appears in the list in the Manage Users and Groups editor.

Creating an Allowed User

In the Manage Users and Groups editor for Active Directory with Local Override, click **New Allowed User**. The **Create New Allowed User** dialog box opens.

Figure 3-10
Create New Allowed User dialog box

The screenshot shows the "Create New Allowed User" dialog box. It features a title bar with standard window controls. The main area contains a "User Name:" text input field. Below it is a section titled "Manage Groups the User belongs to" with two radio button options: "Show all extended groups (Please note it may take a long time to get the list)" and "Show groups starting with:". The second option is selected. To the right of the second option is a text input field for filtering and a "Find" button. Below this are two list boxes: "Available Groups" on the left and "User belongs to Groups" on the right. Between these lists are ">>>>" and "<<<<" buttons. At the bottom of each list box are "Add All" and "Remove All" buttons respectively. At the very bottom of the dialog are "OK", "Cancel", and "Help" buttons.

User Name. The name is not case-sensitive and can contain spaces. Note that it is not possible to verify that the user actually exists in the remote directory, and an incorrectly entered user name will never authenticate to the system.

Show all extended groups. Returns a list of all extended groups.

Show groups starting with. Filters the list of available groups according to the string entered. Use this field to refine the list of available groups.

Available groups. Lists the recognized groups to which the user can be assigned.

User belongs to groups. Lists the groups to which the user is currently assigned.

Add all. Associates all groups with the user.

Remove all. Disassociates all displayed groups from the user.

Note: An allowed user can be associated with extended groups only if extended groups are enabled for Active Directory with Local Override. If extended groups are not enabled, user selection fields are not displayed.

Creating an allowed user requires the user name to be specified. The user can also be associated with groups.

1. In the Create New User dialog box, specify the user name.
2. If necessary, associate the user with extended groups.
3. Click OK. The new allowed user appears in the list in the Manage Users and Groups editor.

Roles

Roles Overview

Roles provide a way to manage user and group access to system functionality. Roles are assigned to users and groups and work in conjunction with a security provider.

Each role created has associated actions that represent the permissions and level of control that the user or group assigned to the role has. For example, a basic user role can be created. The basic user role is assigned a limited set of actions for access to the system and the ability to view the contents of the repository. The basic user role does not have the associated actions to define servers, add other users, or define system configurations that would impact other users and groups.

However, an advanced user role is needed to perform administrative tasks, such as deleting users, creating groups, and defining additional roles. In this case, a less restricted role can be created with more control over the application domain and assigned to a very small set of users.

The list of available actions are defined within the system and cannot be edited by the user assigning them.

If the user belongs to several groups, the roles assigned to that user—an action set—consist of all roles explicitly assigned to the user as well as all roles indirectly assigned through group membership. If the user or group is assigned to several roles, the user or group's action set consists of all roles explicitly assigned as well as all roles indirectly assigned through group membership. Users and groups must be managed per security provider; roles are managed across security providers. For information about managing users and groups, see [Chapter 3](#).

Use the Server Administration tool of Deployment Manager to manage role definitions and to modify the users and groups assigned to roles.

Actions

A role consists of a list of actions. These actions are defined by the system and cannot be changed.

PASW Collaboration and Deployment Services actions

- **Access Contents and Folders.** Access the Content Repository.
- **Configuration.** Modify repository settings.
- **Configure Model.** Configure models for scoring.
- **Create Subscriptions.** Create individual subscriptions to repository objects, such as folders, files, jobs, etc. The subscribers receive e-mail messages when changes are made to the corresponding objects.
- **Define and Manage Notifications.** Define and manage notifications for multiple individuals for events such as job success or failure.

- **Define Credentials.** Create, view, and modify security credentials for execution servers.
- **Define Custom Properties.** Define and modify custom properties for objects within repository.
- **Define Datasources.** Define and modify data sources.
- **Define Message Domains.** Define and modify domains for JMS messaging.
- **Define Server Clusters.** Define and modify execution server clusters.
- **Define Servers.** Define and modify execution servers.
- **Define Topics.** Define and modify topic hierarchy for repository.
- **Job Edit.** Create and modify jobs. Note that job visibility to a user is determined by permissions.
- **Job Run.** Execute jobs. Note that job visibility to a user is determined by permissions.
- **Manage Locks** Manage locks that users create on repository resources, for example, unlock resources locked by others.
- **Manage Enterprise View.** Create, modify, and delete Enterprise Views, Application Views, and Data Provider Definitions.
- **Manage Subscriptions.** Manage other users' subscriptions, and delete subscriptions.
- **MIME Types.** Manage MIME type mappings for repository.
- **Repository Index.** Reindex the contents of repository.
- **Run Custom Dialogs** Run PASW Statistics custom dialogs.
- **Run Report Dynamically.** Run dynamic reports, such as Business Intelligence Reporting Tools (BIRT) reports, in Deployment Portal.
- **Schedules.** Manage job schedules.
- **Score Model.** Score models.
- **Show All Versions.** View all versions of objects (labeled and unlabeled) in Deployment Portal. By default, users are able to see only labeled versions in Deployment Portal.
- **Show latest.** View only the latest version of objects.
- **Submit Work** Submit work (for example, reports) for processing by PASW Collaboration and Deployment Services.
- **User Preference Administration.** Manage the preferences of other users.
- **View Expired Files.** View expired content, such as files and jobs.
- **View Model Management Dashboard.** View model management dashboards in Deployment Manager and Deployment Portal.

Note: Show latest action is a subset of Show All Versions and if a user has both actions, Show All Versions supersedes Show latest.

Administrators Role

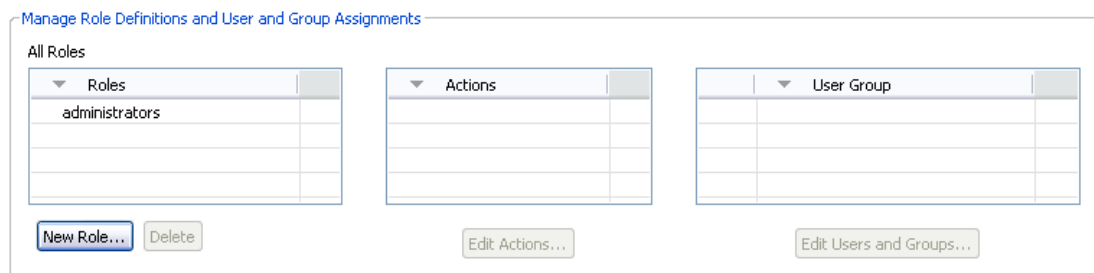
The system includes a predefined *administrators* role that cannot be modified. This role is associated with all actions available in the system. Any user assigned to this role will be able to perform any action in the system. In addition, some functionality not controlled by actions, such as export and import of repository content, is available only to users assigned to this role.

Due to the breadth of control available to administrators, care should be exercised when assigning users to this role. Assign only those users who need access to all functionality in the system. Users who need only a subset of actions should be assigned to custom roles. For more information, see the topic [Creating a New Role](#) on p. 38.

Manage Role Definitions

To work with roles, choose Server Administration from the Tools menu, select an repository Server, and log in. Double-click the Roles icon for the server to access the Manage Role Definitions editor.

Figure 4-1
Manage Role Definitions and User and Group Assignments



Please note that all changes made in this editor will take immediate effect and there will be no way to cancel modifications.

All roles. Provides a list of all roles available for the security provider. When new roles are added, this list is populated with entries. To add a new role to the system, click the New Role button. To delete a role, select the role and click the Delete button. Select a role from this list to view its associated actions.

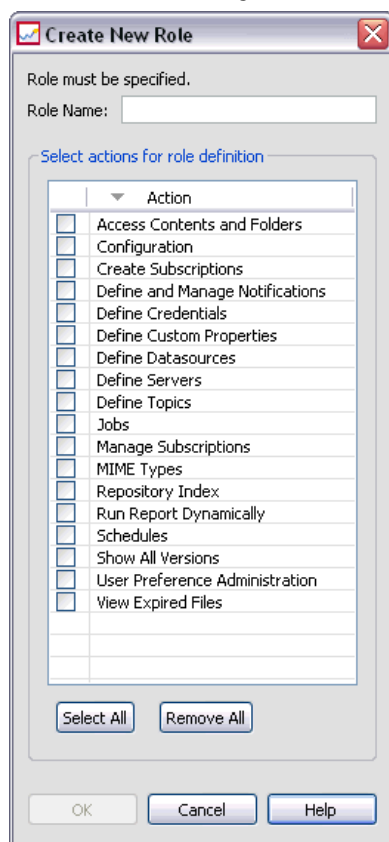
Definition of roles. Provides a list of actions associated with a selected role. To edit the actions associated with a selected role, click the Edit Actions button.

Users and Groups Assigned to Role. A list of the users and groups assigned to a selected role. To edit the users and groups list for a selected role, click the Edit Users and Groups button.

Creating a New Role

To create a new role, click the New Role button in the Roles editor. The Create New Role dialog box opens.

Figure 4-2
Create new role dialog



Role Name. A text string to identify the role. The role name should be unique and not duplicate another role name.

Action. Contains all actions defined and available within the system. Initially, a role has no actions associated with it.

Select the box next to an action to assign it to the role. Alternatively, click the **Select All** button to add all actions to the role. Clicking the **Remove All** button clears all actions from the role. The list of actions can be sorted by clicking on the **Action** column. Click **OK** to create and save the role.

Editing a Role

To edit the list of actions assigned to a role, select the role to edit in the Roles editor and click the **Edit Actions** button. The **Edit Role** dialog box opens.

Role name. A text string to identify the role. The role name should be unique and not duplicate another role name.


Action. Contains all actions defined and available within the system. Initially, a role has no actions associated with it.

Select the box next to an action to assign it to the role. Alternatively, click the **Select All** button to add all actions to the role. Clicking the **Remove All** button clears all actions from the role. The list of actions can be sorted by clicking on the **Action** column. Click **OK** to save the modified role definition.

Editing Users and Groups Assigned to a Role

Once roles have been defined, the roles need to be associated with users and groups to define levels of access. To assign users and groups to a role, from the Roles editor, click the **Edit Users and Groups** button. The **Edit Users and Groups for Role** dialog box opens.

Figure 4-3
Edit Users and Groups for Role dialog box



Two options exist for viewing users and groups that can be assigned to roles:

- **Show all available users/groups.** Provides a list of all users and groups available for all security providers.
- **Shows users/groups starting with.** Filters the available list of users and groups according to the search options.


The Available Users/Groups list is populated with users and groups according to the search option. Select a user or group and click the >>>> button to assign it to the role. To remove a user or group from a role, select the user or group in the Users/Groups Assigned to Role list and click the <<<< button. When finished, click **OK**.

Removing a Role

To remove a role:

1. From the Roles editor, select the role to remove.
2. Click the Delete button. A confirmation dialog box opens.

Figure 4-4
Confirm Delete dialog box



3. Click OK to verify that the role should be removed.

The role is removed from the system.

Security Providers

A security provider is responsible for verifying the credentials supplied by a user against a particular user directory. PASW Collaboration and Deployment Services includes an internal directory for authentication, but an existing enterprise user directory can also be used. Available providers include:

- **Native (or local user repository).** The internal security provider for PASW Collaboration and Deployment Services, in which users, groups, and roles can all be defined. The native provider is always active and cannot be disabled.
- **Active Directory®.** The Microsoft version of Lightweight Directory Access Protocol (LDAP) for authentication, authorization, and security policies. Users and groups for this provider must be defined directly in the Active Directory framework. After configuring Active Directory for use with PASW Collaboration and Deployment Services, the system can authenticate a user against the Active Directory server while maintaining the permissions and access rights associated with that user. In contrast to the native provider, this provider can be enabled or disabled. For additional information about Active Directory, see the original vendor's documentation.
- **Active Directory with local override.** A provider that leverages Active Directory but allows the creation of extended groups and allowed-users filters. An extended group contains a list of users from Active Directory but exists outside of the Active Directory framework. An allowed-users filter restricts the list of Active Directory users that can authenticate against the system to a defined set. This provider can be enabled or disabled.
- **IBM i.** IBM i user profiles directory can be used to authenticate PASW Collaboration and Deployment Services users. This provider can be enabled or disabled. If IBM i security provider is used with single sign-on-enabled PASW Collaboration and Deployment Services installation, EIM (Enterprise Identify Management) must be enabled. Additionally, */QIBM/UserData/Java400/ext/eim.jar* must be copied into the library directory of the PASW Collaboration and Deployment Services application server if the application server is running on a non-IBM i host.
- **OpenLDAP®.** An open-source LDAP implementation for authentication, authorization, and security policies. Users and groups for this provider must be defined directly using LDAP tools. After configuring OpenLDAP for use with PASW Collaboration and Deployment Services, the system can authenticate a user against the OpenLDAP server while maintaining the permissions and access rights associated with that user. This provider can be enabled or disabled.
- **SiteMinder®.** Computer Associates enterprise-scale Web applications access management system for policy-based authentication and authorization. Users and groups for this provider are defined in an external LDAP-based directory, whereas access policies are defined in SiteMinder. The provider also allows the creation of extended groups and allowed-users filters.

Important! SiteMinder security provider is not enabled by default. To enable SiteMinder authentication, <PASW Collaboration and Deployment Services installation directory>/paswa4.0/optional/siteminder.package must be installed.

Security Providers in Deployment Manager

Before performing any actions with security providers, navigate to the administrative interface that controls this functionality.

1. From the Tools menu, choose Server Administration.
2. On the Server Administration tab, log in to a PASW Collaboration and Deployment Services server.
3. Double-click the Configuration icon for the server to expand the hierarchy.
4. Double-click the Security Providers icon to expand the hierarchy.
5. Double-click the type of security provider to configure.

Figure 5-1
Accessing security providers



Configuring Security Providers

Each type of security provider has settings specific to the type of authentication and authorization system being used.

Native

The native security provider is internal to PASW Collaboration and Deployment Services and does not contain any settings to configure.

Active Directory

To configure Active Directory settings, double-click the Active Directory icon.

Figure 5-2
Active Directory configuration

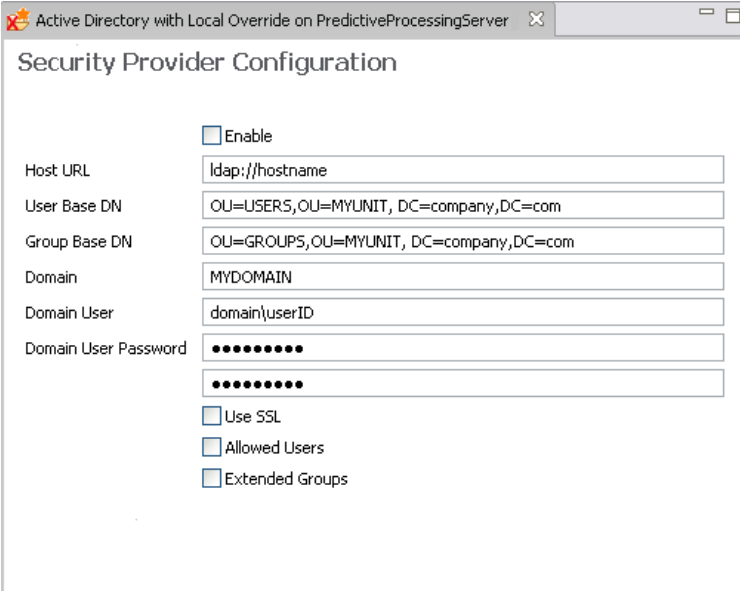
- **Enable.** Enables and disables the use of **Active Directory** as a service provider.
- **Host URL.** URL for the Active Directory server. The default port for LDAP is 389.
- **User Base DN.** Base **distinguished name** for user searches.
- **Group Base DN.** Base distinguished name for group searches.
- **Domain.** The DNS namespace to which the user is logging in.
- **Domain User.** A user ID to perform searches, specified in the format *domain\username*. The specified name must have the proper permissions to look up and authenticate users.
- **Domain User Password.** For security, the specified domain user password appears in a hashed asterisk (*) format. Type the value in both password fields to verify the correct value.
- **Use SSL.** Select to use secure sockets for communication with the Active Directory server.

After you have modified the settings, choose Save from the File menu.

Active Directory with Local Override

To configure the settings for the Active Directory with local override security provider, double-click the Active Directory with Local Override icon. Most of the settings are identical to the Active Directory settings. However, local override offers two additional settings.

Figure 5-3
Active Directory with Local Override configuration



The screenshot shows a window titled "Active Directory with Local Override on PredictiveProcessingServer" with a sub-header "Security Provider Configuration". The configuration includes the following fields and options:

- Enable
- Host URL: ldap://hostname
- User Base DN: OU=USERS,OU=MYUNIT, DC=company,DC=com
- Group Base DN: OU=GROUPS,OU=MYUNIT, DC=company,DC=com
- Domain: MYDOMAIN
- Domain User: domain\userID
- Domain User Password: (two rows of masked characters)
- Use SSL
- Allowed Users
- Extended Groups

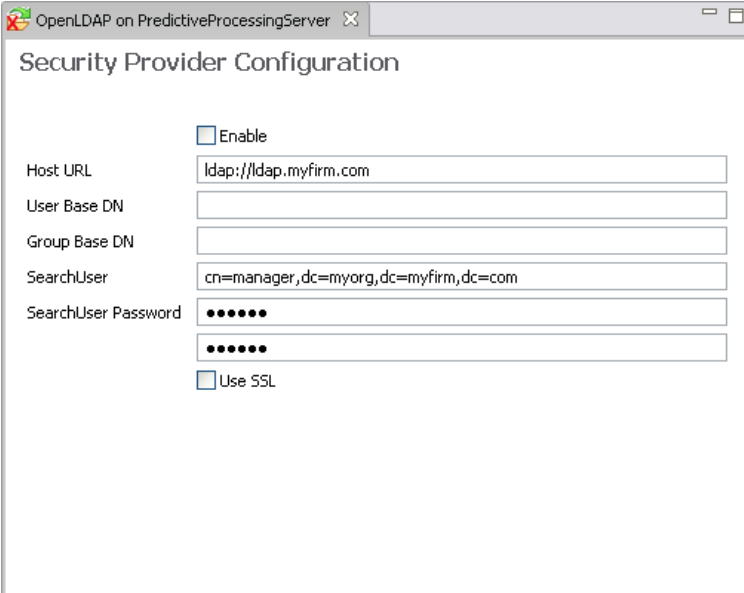
- **Allowed Users.** Enables and disables the use of allowed users, which allows only users on a locally defined list to be authenticated in Active Directory.
- **Extended Groups.** Enables and disables the use of extended groups, which allow a group of Active Directory users to be defined. Active Directory users can be assigned to these local groups.

After you have modified the settings, choose **Save** from the **File** menu.

OpenLDAP

To configure OpenLDAP settings, double-click the OpenLDAP icon.

Figure 5-4
OpenLDAP configuration



OpenLDAP on PredictiveProcessingServer

Security Provider Configuration

Enable

Host URL:

User Base DN:

Group Base DN:

SearchUser:

SearchUser Password:

Use SSL

- **Enable.** Enables and disables the use of OpenLDAP as a security provider.
- **Host URL.** The path to the LDAP server, usually a DNS resolvable name or an IP address (for example, *ldap://*). The default port for LDAP is 389.
- **User Base DN.** Base distinguished name for user searches.
- **Group Base DN.** Base distinguished name for group searches.
- **SearchUser.** A user ID to perform searches, specified in a distinguished name format. The specified name must have the proper permissions to look up and authenticate users.
- **SearchUser Password.** For security, the specified domain user password appears in a hashed asterisk (*) format. Type the value in both password fields to verify the correct value.

After you have modified the settings, choose **Save** from the File menu.

IBM i

To configure IBM i security provider, double-click the IBM i icon.

Figure 5-5
IBM i configuration

- **Enable.** Enables and disables the use of an IBM i system as a security provider.
- **IBM i Server.** The path to the IBM i system, usually a DNS resolvable name or an IP address. If you are using IBM i security provider with Enterprise Identity Management (EIM) to enable single sign-on to PASW Collaboration and Deployment Services, then this value must match the EIM target registry value. If the EIM target registry value is a fully qualified name of the host, enter fully qualified host name.
- **User Profile.** The user profile used to perform directory searches on the IBM i system.
- **Password.** The password for the specified IBM i profile. For security, the specified domain user password appears in a hashed asterisk (*) format. Type the value in both password fields to verify the correct value.
- **Enable EIM Lookup.** For single sign-on enabled PASW Collaboration and Deployment Services installations, indicates that Enterprise Identity Management is enabled.
- **EIM Server.** Enterprise Identity Management host URL.
- **EIM User.** The user name for Enterprise Identity Management host login.
- **EIM Password.** The password for the specified Enterprise Identity Management user.

Note: Any IBM i user profile can be used for directory searches, but the list profiles that is returned will be filtered based on the authority of the profile used for the search. Specifying a QSECOFR level user will return all profiles on the system. Using a user with fewer privileges will result in fewer profiles being returned based on the user profiles security settings.

IBM i User and Group Permissions

If an IBM i user profile is intended to be used as a group, other IBM i profiles must be assigned to the profile before it is assigned PASW Collaboration and Deployment Services permissions. Otherwise, the permissions are not inherited by other IBM i users. For example, if an IBM i user *test* is created, assigned permissions in PASW Collaboration and Deployment Services, and

then assigned as a group to IBM i user *test2*, *test2* does not inherit the previous permissions of *test* in PASW Collaboration and Deployment Services. However, if *test2* is associated with *test* before PASW Collaboration and Deployment Services permissions of *test* are defined, *test2* does inherit the permissions.

SiteMinder Security Provider

To configure SiteMinder settings, double-click the SiteMinder icon.

Figure 5-6
SiteMinder configuration

The screenshot shows the 'Security Provider Configuration' window. At the top, there is a checkbox labeled 'Enable'. Below it are several input fields with the following values:

Local Computer	127.0.0.1
Site Minder IP	127.0.0.1
Agent Secret
Directory Root	OU=NEWYORK, DC=spss, DC=COM
Organization Root	OU=USERS, OU=NEWYORK, DC=spss, DC=COM
Directory Server IP	127.0.0.1
Directory User Name	SiteMinder
Directory User Password
Directory Name	SPSS AD
Directory Search User Query	(&(objectCategory=user)(sAMAccountName=\$ID))
Directory Search Group Query	(&(objectCategory=group)(sAMAccountName=\$ID))

- **Enable.** Enables and disables the use of SiteMinder as a security provider.
- **Local Computer:** PASW Collaboration and Deployment Services host address or domain name.
- **SiteMinder IP.** SiteMinder policy server address or domain name.
- **Agent secret.** Agent secret for agent *spss-agent*, used for SiteMinder policy server login.
- **Directory Root.** Directory root base DN, for example DC=domain name part 1, DC=domain name part 2
- **Organization Root.** Organization root, for example OU=unit, DC=domain name part 1, DC=domain name part 2 ...
- **Directory Server.** Directory server address or domain name.
- **Directory User Name.** The name of the user for directory login.
- **Directory User Password.** The password for directory login.
- **Directory Name.** LDAP directory name.
- **Directory Search User Query.** LDAP query for user-based authentication.
- **Directory Search Group Query.** LDAP query for group-based authentication.

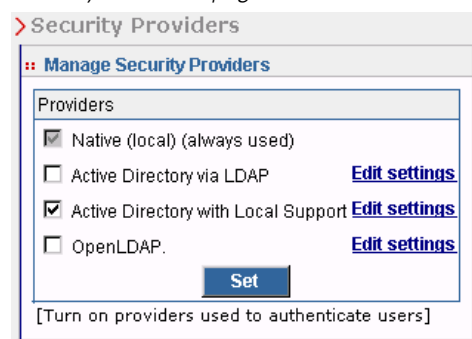
After you have modified the settings, choose Save from the File menu.

Security Providers in the Browser-based Deployment Manager

To access the Security Providers page:

- ▶ Click Security Providers in the navigation list. The Security Providers page appears.

Figure 5-7
Security Providers page



To modify the security providers used:

- ▶ Select (to enable) or deselect (to disable) the check boxes next to the security provider.
- ▶ Click Set.

Configuring Security Providers

Each type of security provider has settings specific to the type of authentication and authorization system being used.

Native

The native (or local) security provider is inherent to the system and cannot be removed. Users can be added to the native security system, but it cannot be disabled or configured.

Active Directory

To configure Active Directory settings, click Edit settings to the right of the Active Directory via LDAP check box. The current settings appear. You can modify the following settings:

- **Host URL.** URL for the **Active Directory** server. The default port for LDAP is 389.
- **User Base DN.** Base **distinguished name** for user searches.
- **Group Base DN.** Base distinguished name for group searches.
- **Domain.** The DNS namespace to which the user is logging in.
- **Domain User.** A user ID to perform searches, specified in the format *domain\username*. The specified name must have the proper permissions to look up and authenticate users.

- **Domain User Password.** For security, the specified domain user password appears in a hashed asterisk (*) format.
- **Use SSL.** Specify true to use secure sockets for communication with the Active Directory server.

After you have modified the settings, click Set.

Active Directory with Local Override

The Active Directory with local override security provider option allows Active Directory to be used with the additional options of a local principal filter and the ability to specify local groups. To configure Active Directory with local override settings, click Edit settings to the right of the Active Directory with Local Override check box. The current settings appear. Most of the settings correspond to those for Active Directory. However, the following two options are also available:

- **Allowed Users.** Enables (true) and disables (false) the use of allowed users, which allows only users on a locally defined list to be authenticated in Active Directory.
- **Extended Groups.** Enables (true) and disables (false) the use of extended groups, which allow a group of Active Directory users to be defined. Active Directory users can be assigned to these local groups.

After you have modified the settings, click Set.

IBM i

When repository is installed on IBM i, the IBM i user profile directory will be used to authenticate repository logins. To configure IBM i settings, click Edit settings to the right of the IBM i check box. The current settings appear. You can modify the following settings:

- **IBM i Server.** The path to the IBM i system, usually a DNS resolvable name or an IP address. If you are using IBM i security provider with Enterprise Identity Management (EIM) to enable single sign-on to PASW Collaboration and Deployment Services, then this value must match the EIM target registry value. If the EIM target registry value is a fully qualified name of the host, enter fully qualified host name.
- **User Profile.** The user profile to used to perform directory searches on the IBM i system.
- **Password.** The password for the specified IBM i profile. For security, the specified domain user password appears in a hashed asterisk (*) format.
- **Enable EIM Lookup.** For single sign-on enabled PASW Collaboration and Deployment Services installations, the value of true indicates that Enterprise Identity Management is enabled.
- **EIM Server.** Enterprise Identity Management host URL.
- **EIM User.** The user name for Enterprise Identity Management host login.
- **EIM Password.** The password for the specified Enterprise Identity Management user.

Note: Any IBM i user profile can be used for directory searches, but the list profiles that is returned will be filtered based on the authority of the profile used for the search. Specifying a QSECOFR level user will return all profiles on the system. Using a user with fewer privileges will result in fewer profiles being returned based on the user profiles security settings.

OpenLDAP

To configure OpenLDAP settings, click **Edit settings** to the right of the OpenLDAP check box. The current settings appear. You can modify the following settings:

- **Host URL.** The path to the LDAP server, usually a DNS resolvable name or an IP address (for example, *ldap://*). The default port for LDAP is 389.
- **User Base DN.** Base distinguished name for user searches.
- **Group Base DN.** Base distinguished name for group searches.
- **SearchUser.** A user ID to perform searches, specified in a distinguished name format. The specified name must have the proper permissions to look up and authenticate users.
- **SearchUser Password.** For security, the specified domain user password appears in a hashed asterisk (*) format.

After you have modified the settings, click **Set**.

SiteMinder

To configure SiteMinder settings, click **Edit settings** to the right of the SiteMinder check box. The current settings appear. You can modify the following settings:

- **Local Computer:** PASW Collaboration and Deployment Services host address or domain name.
- **SiteMinder IP.** SiteMinder policy server address or domain name.
- **Agent secret.** Agent secret for agent *spss-agent*, used for SiteMinder policy server login.
- **Directory Root.** Directory root base DN, for example DC=domain name part 1, DC=domain name part 2
- **Organization Root.** Organization root, for example OU=unit, DC=domain name part 1, DC=domain name part 2 ...
- **Directory Server.** Directory server address or domain name.
- **Directory User Name.** The name of the user for directory login.
- **Directory User Password.** The password for directory login.
- **Directory Name.** LDAP directory name.
- **Directory Search User Query.** LDAP query for user-based authentication.
- **Directory Search Group Query.** LDAP query for group-based authentication.

After you have modified the settings, click **Set**.

Single Sign-On

Single sign-on (SSO) is a method of access control that enables a user to log in once and gain access to the resources of multiple software systems without being prompted to log in again. PASW Collaboration and Deployment Services provides the single sign-on capability by initially authenticating users through an external directory service based on **Kerberos** security protocol, and subsequently using the credentials in all PASW Collaboration and Deployment Services applications, for example, Deployment Manager, Deployment Portal, or a portal server without additional authentication.

Single sign-on configuration is performed on the Server Administration tab of Deployment Manager. Note that a number of prerequisites must be in place before single sign-on can be enabled. For more information, see PASW Collaboration and Deployment Services installation and configuration documentation.

Configuring Single Sign-on

- ▶ Choose Server Administration from the Tools menu, log into a PASW Collaboration and Deployment Services server, and double-click the Single Sign-On icon. Single Sign-on Provider editor opens.

Figure 6-1
Single Sign-On Provider Editor

Single Sign On Provider Configuration	
<input type="checkbox"/> Enable	
Security Provider	Active Directory
KDC Host Address	kdc.mycompany.com
Kerberos Realm	MYCOMPANY.COM
Host Address	paswserver.mycompany.com
Kerberos Service Principal	HTTP/paswserver.mycompany.com@MYCOMPANY.COM
Kerberos Service Principal Password

Kerberos Key Table URL	FILE:C:/keytab/krb5.keytab
JAAS Configuration File	#USE_SUPPLIED#

- **Enable.** Enables or disables the use of single sign-on provider.
- **Security provider** A configured external security providers, such as Windows Active Directory. Local security provider cannot be selected.

- **Kerberos Key Distribution Center Host Address** Fully qualified name of the Kerberos Domain controller host. For Windows Active Directory, this is the name of the host where Microsoft Active Directory Services are installed.
- **Kerberos Realm** The Kerberos realm. For Active Directory, this is the domain name.
- **PASW Host** The address of PASW Collaboration and Deployment Services repository host, for example, `http://paswhost.mycompany.com:8080`.
- **Kerberos Service Principal Name** The user name for the Kerberos Service Principal.
- **Kerberos Service Principal Password** The password of the user Kerberos Service Principal.
- **Kerberos Key Table URL** The URL of the keytab file for Kerberos principals authentication.
- **JAAS Configuration File** The path of JAAS (Java Authentication and Authorization Service) configuration file on the PASW Collaboration and Deployment Services host file system. If specified, it overrides the default JAAS configuration. Depending on the application server, this may be necessary to configure the JRE to support SSO.

JAAS Configuration

At a minimum, at least one JAAS configuration for JGSSServer must be provided with the following parameters:

- **JGSSServer** required
 - **KerberosLocalUser** optional
 - **JDBC_DRIVER_01** optional
- For Sun JRE, the following default JGSSServer configuration is created:

```
JGSSServer {
  com.sun.security.auth.module.Krb5LoginModule required
  storeKey="true"
  doNotPrompt="true"
  realm=<realm name>
  useKeyTab="true"
  principal=<name>
  keyTab=<path>
  debug=false;
};
```

- Optional KerberosLocalUser configuration is used to allow NTLM bypass. This configuration allows the user to create a Kerberos credential when client browser sends a NTLM token (instead of a Kerberos token) during the negotiation challenge. Note that on Windows system, browsers on the same machine, where PASW Collaboration and Deployment Services server is installed, will always send NTLM token. All NTLM requests to PASW Collaboration and Deployment Services may be disabled by omitting this configuration from their JAAS configuration file.

```
KerberosLocalUser {
  com.sun.security.auth.module.Krb5LoginModule required
  useTicketCache="true"
  debug=false;
};
```

- ▶ Optional JDBC_DRIVER_01 configuration is used for Kerberos authentication to database servers.

```
JDBC_DRIVER_01 {  
  com.sun.security.auth.module.Krb5LoginModule required  
  useTicketCache="true"  
  debug=false;  
};
```

- ▶ It is also possible to specify appropriate login module class name, requirement type, and other options that the login module requires for each JAAS configuration. The login module class must be in class path. For more information, see JRE and application server vendor documentation.


Repository Configuration

PASW Collaboration and Deployment Services provides a number of options for configuring its components, ranging from the templates that are used for the user interface to the messages that appear on the Login screen.

To access any of these options, in the browser-based Deployment Manager:

- ▶ Click Configuration in the navigation list. The Configuration page opens.

Figure 7-1
Configuration page



- ▶ In the Configuration list, click the link that corresponds to the property that you want to configure.

Each property configuration screen has two buttons, Set and Use Default. Once a configuration is made, click the Set button for the new setting to take effect. To restore a value to the original system configuration, click the Use Default button.

Note: Certain configuration options listed below are intended for optional PASW Collaboration and Deployment Services components or other SPSS Inc. products, such as PASW Statistics or ShowCase. The options are not available if the components are not installed.

Administrator

The Administrator configuration option allows you to specify the location of the templates used to generate the administrative user interfaces. By default, the system uses the path established by the installation program.

To modify the templates directory:

- ▶ In the Configuration list, under Administrator click Templates. The current templates directory appears in the Templates text box.
- ▶ In the Templates text box, enter the new path of the directory that contains the templates that you want to use.
- ▶ Click Set. The path that you specified becomes the default path for the system to access templates.
- ▶ To return to the system-defined default, click Use Default. This option restores the default directory that was established when you installed the system.

Cache Provider

The Cache Provider option allows you to specify and configure the data cache provider class. By default, Ehcache (*com.spss.cache.service.ehcache.EhcacheProvider*) is used. In clustered PASW Collaboration and Deployment Services installations, additional options allow configuring Ehcache for automatic discovery of peers participating in a cluster using a multicast group.

Alternatively, Oracle Coherence can be used as repository cache. To enable Oracle Coherence:

- ▶ Obtain and license Coherence components from Oracle. Coherence*. *jar* files must be placed into the *<PASW Collaboration and Deployment Services installation location>/components/cache-provider*.
- ▶ Install the *<PASW Collaboration and Deployment Services installation location>/pasw4.0/optional/coherence_cache_provider.package*.
- ▶ Specify *com.spss.cache.service.coherence.CoherenceCacheProvider* as a cache provider in configuration settings.

To modify the settings, click the corresponding option under Cache Provider in the Configuration list. See the following table for link names, descriptions, and valid settings.

Name	Description	Settings
Cache Provider Class Name	Cache adapter class name.	Class name.
Multicast Group Address	For Ehcache, the multicast group address.	Valid network address.
Multicast Group Port	For Ehcache, a dedicated port for the multicast heartbeat traffic.	Valid port number.
Override Default Values	For Ehcache, if the option is enabled, the provider will use <i>Multicast Group Address</i> and <i>Multicast Group Port</i> values to override the defaults.	Unselected by default.

Custom Dialogs

PASW Statistics custom dialogs configuration options allow you to specify the default PASW Statistics server for executing custom dialogs and the credentials used to connect to the server.

To modify the settings, click the corresponding option under Notification in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-1
Notification configuration options

Name	Description	Settings
PASW Statistics server	The repository name or URI of PASW Statistics server used to execute custom dialog syntax. Alternatively, the name or URI of a server cluster can be specified. In that case, a server will automatically be selected from the cluster based on availability. If no server is specified, the default server will be selected by using an available server from the first valid server cluster definition that is found. If no valid clusters are found, the first valid server that is found will be used.	A string value corresponding to the repository name or URI of the server object, for example <code>spsscr:///id=0a30063bc975ede400</code> . The URI can be found in the object properties. For more information, see Deployment Manager documentation.
PASW Statistics server credential	The credential used to connect to the PASW Statistics server when executing custom dialog syntax.	A string value corresponding to the repository name or URI of the credential object.

Data Service

Data Service configuration options allow you to specify parameters for optimizing Data Service connections.

To modify the settings, click the corresponding option under Deployment Portal in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-2
Data Service configuration options

Name	Description	Settings
Active Connectors Maximum Number	Maximum number of active connections.	Integer value. Default is 5.
Idle Connectors Maximum Number	Maximum number of idle connections.	Integer value. Default is 5.

Deployment Manager

The Deployment Manager configuration option allows you to specify the protocol timeout for communication between Deployment Manager and repository. Specify the number of seconds the Deployment Manager client should wait for an repository server. Use a larger value if timeout errors are received for server transactions.

To modify the protocol timeout:

- ▶ In the Configuration list, under Deployment Manager click Protocol Timeout. The current value appears.
- ▶ In the Protocol Timeout text box, enter the desired number of seconds.
- ▶ Click Set. The value you specified becomes the timeout value.
- ▶ To return to the system-defined default, click Use Default. This option restores the default value that was established when you installed the system.

Deployment Portal

Deployment Portal configuration options allow you to specify authentication settings and the report timeout limit for the web-based Deployment Portal application.

To modify the settings, click the corresponding option under Deployment Portal in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-3
Deployment Portal configuration options

Name	Description	Settings
Configured Authentication Criteria Class	The Java class name used to provide authentication information for the Deployment Portal application. Defaults to <i>com.spss.er.internal.configuration.ConfiguredAuthenticationImpl</i> and is set in the classpath of the application server. The class must conform to the authentication criteria interface provided by Deployment Portal (<i>com.spss.er.internal.configuration.ConfiguredAuthenticationInterface.java</i>).	Class name.
Use Configured Authentication Criteria	Allows user to pass authentication information to Deployment Portal using the Configured Authentication Criteria, hence bypassing the Login screen.	Unselected by default.

Deployment Portal Scoring

The Batch Scoring Row Limit configuration option allows you to specify the maximum number of rows that may be batch scored from a selected data set.

To modify the row limit:

- ▶ In the Configuration list, under Deployment Portal Scoring, click Batch Scoring Row Limit. The current value appears.
- ▶ In the Batch Scoring Row Limit text box, enter the desired number of minutes.

- ▶ Click Set. The value you specified becomes the timeout value.
- ▶ To return to the system-defined default, click Use Default. This option restores the default value that was established when you installed the system.

Enterprise View

Enterprise View configuration options allow you to specify settings for working with a PASW Statistics Data File Server. To modify the settings, click the corresponding option under Enterprise View in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-4
Enterprise View configuration options

Name	Description	Settings
Maximum CQL query columns	The maximum number of rows returned by CQL (Common Query Language) queries.	Integer value. Default is 2.
PASW Statistics Data File Additional Servers	This setting is used to specify additional PASW Statistics data file servers that can be used to retrieve metadata from PASW Statistics data files. .	A semicolon delimited list of host:port values, for example, server2:18886;server3:18886
PASW Statistics Data File Load Balance	The load balancing setting controls whether multiple PASW Statistics data file servers are used in failover mode or load balancing mode when retrieving metadata from PASW Statistics data files. In failover mode, the list servers are used in sequential order. If the first does not work, the second is used, etc. When load balancing is turned on, one of the available servers is selected at random. This setting has no effect unless additional PASW Statistics Data File Additional Servers are specified.	Selected by default.
PASW Statistics Data File Server Host	The name of the PASW Statistics Data File Server used for accessing PASW Statistics data files. If a host is not specified, the localhost will be used.	Any valid IP address or hostname.
PASW Statistics Data File Server Port	The port for the PASW Statistics Data File Server. If the port is not specified, the default port will be used.	A valid port number.
PASW Statistics Data File Server Secure	Indicator of whether or not SSL should be used when communicating with the PASW Statistics Data File Server. The default value of false means secure sockets are not used.	True or false. Default is false.

Help

The Help configuration option allows you to specify the location of the Help system for Deployment Manager. By default, the system uses the path that was established by the installation program.

To modify the Help directory:

- ▶ In the Configuration list, under Help click Directory. The current Help directory appears in the Help text box.
- ▶ In the Help text box, enter the path of the directory that contains the Help system.
- ▶ Click Set. The path that you specified becomes the default path for the system to access online help.
- ▶ To return to the system-defined default, click Use Default. This option restores the default directory that was established when you installed the system.

Notification

Notification configuration options allow you to specify SMTP mail settings and enable notification service performance tuning. For more information, see the topic [Optimizing Notification Service Performance](#) in Chapter 10 on p. 87.

To modify the settings, click the corresponding option under Notification in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-5
Notification configuration options

Name	Description	Settings
Binary Content Enabled	Enables binary content, such as e-mail attachments, for notification messages.	Selected by default.
Core Event Collector Pool Size	The number of threads to keep in the event collector pool, even if they are idle.	Integer value. Default is 16.
Distinct Recipients	If the check box is selected, notification messages will be generated only for unique recipients. Otherwise, duplicate addresses will not be removed, and the recipients will get the messages generated by all of their individual subscriptions and notifications that match the given notification event. The option should be changed only for debugging purposes.	Selected by default.
Event Collector Enabled	Defines whether notification events should be processed by the service.	Selected by default.

Name	Description	Settings
Event Collector Pool Keep Alive Time	When the number of threads is greater than the core number of threads in the event collector pool, this is the maximum time in seconds that excess idle threads will wait for new events before terminating.	Integer value. Default is 32.
Event Inheritance Enabled	Defines whether derived notification events should be processed by the service.	Unselected by default.
Event Noise Filter	Filter out notification events that do not have matching subscriptions with associated notification providers or subscribers early in the process.	True or false. Default is true.
Event Noise Filter Cache	Defines a maximum size of the LRU cache to use during event noise filtering.	Integer value. Default is 2048.
Event Noise Filter String Keys	Use strings instead of hash codes to identify notification events.	Unselected by default.
Event Queue Storage Commit Batch Size	Sets the commit batch size for the persistent storage for the incoming notification events. The notification service should be restarted for the changes to take effect.	Integer value. Default is 32.
Maximum Event Collector Pool Size	The maximum number of threads allowed in the event collector pool.	Integer value. Default is 32.
Message Bus Enabled	Defines whether notification messages should be sent to the JMS Message Bus.	Selected by default.
Message Bus Filter Enabled	Defines whether only the notifications of interest should be sent to the JMS Message Bus.	Selected by default.
Notification Auditor Enabled	Defines whether notification service should interface with the auditing service.	Selected by default.
Notification Queue	Queues incoming notification events until they can be processed by background threads.	True or false. Default is true.
Persistent Event Queue Enabled	Defines whether incoming notification events should be temporarily kept in the persistent storage on the disk to minimize amount of the consumed memory. The notification service should be restarted for the changes to take effect.	Unselected by default.
Persistent Event Queue Size	Defines the maximum size of persistent storage for the incoming notification events (in megabytes). The notification service should be restarted for the changes to take effect.	Integer value. Default is 8 MB.

Name	Description	Settings
Persistent Event Queue Type	Defines storage type for persistent event queue. The notification service should be restarted for the changes to take effect.	Either DISK or JMS. Default is DISK.
Persistent JMS Connection Factory	Defines JNDI name for JMS Connection Factory used to persist incoming notification events. The notification service should be restarted for the changes to take effect.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the JMS Connection Factory.
Persistent JMS Queue	Defines JNDI name for JMS Queue used to persist incoming notification events. The notification service should be restarted for the changes to take effect.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the JMS queue.
Prefer Individual Subscriptions	If the check box is selected, the processing of the subscriptions will take precedence for the users whose individual subscription settings are identical to the settings of notifications created by the administrator. Unselecting the check box will reverse the order of processing.	Selected by default.
SMTP 8 bit MIME	If set to true, and the server supports the 8BITMIME extension, text parts of messages that use the “quoted-printable” or “base64” encodings are converted to use “8bit” encoding if they follow the RFC2045 rules for eight-bit text.	True or false. Default is false.
SMTP Authentication	If true, attempt to authenticate the user using the AUTH command.	True or false. Default is false.
SMTP Connection Timeout	Socket connection timeout value in milliseconds.	Integer value. Default is infinite timeout.
SMTP Distributor Enabled	If the check box is selected, it enables distribution of notification messages via SMTP. repository administrator can disable SMTP distribution to suppress all the e-mails generated by the server. Note that since repository does not store generated e-mail messages, if the SMTP distribution is disabled, all messages will be lost.	Selected by default.
SMTP DSN Notify	The NOTIFY option to the RCPT command for DSN (Delivery Status Notifications, RFC3461).	Either NEVER or some combination of SUCCESS, FAILURE, and DELAY (separated by commas).
SMTP DSN RET	The RET option to the MAIL command for DSN (Delivery Status Notifications, RFC3461).	Either FULL or HDRS.

Name	Description	Settings
SMTP EHLO	If false, do not attempt to sign on with the EHLO command.	True or false. Default is true.
SMTP from e-mail address	The sender or return address to use for notification e-mail.	Any existing SMTP e-mail address.
SMTP Host	The IP address or hostname of the SMTP server used to send mail.	Any valid IP address or hostname.
SMTP Local Host	Local hostname used in the SMTP HELO or EHLO command. Defaults to <i>InetAddress.getLocalHost().getHostName()</i> . Should not normally need to be set if your JDK and your name service are configured properly.	Any valid IP address or hostname.
SMTP Password	Password for SMTP authentication.	Masked password.
SMTP Port	The port used for outgoing mail.	Any valid port number. Default is 25.
SMTP QUIT	If set to true, causes the transport to wait for the response to the QUIT command. If set to false, the QUIT command is sent and the connection is immediately closed.	True or false. Default is false.
SMTP SASL Realm	The SASL (Simple Authentication and Security Layer) realm to use with DIGEST-MD5 authentication.	A deployment-specific and/or server-specific case-sensitive string that identifies the realm or domain from which the principal name should be chosen.
SMTP Send Partial	If set to true, and a message has some valid and some invalid addresses, sends the message anyway, reporting the partial failure with a <i>SendFailedException</i> . If set to false, the message is not sent to any of the recipients if there is an invalid recipient address.	True or false. Default is false.
SMTP Timeout	Socket I/O timeout value in milliseconds.	Integer value. Default is infinite timeout.
SMTP Transfer Protocol	Message transfer protocol.	Either <code>smtp</code> or <code>smtps</code> . Default is <code>smtp</code> while <code>smtps</code> is used to connect to the corresponding service using SSL/TLS.
SMTP Turn on Debug Mode	Toggles debug mode on and off.	True or false. Default is false.
SMTP User	Default user name for SMTP.	Username.
Subscription Identifier Cache	Defines a maximum size of the LRU cache for commonly used subscription identifiers.	Integer value. Default is 2048.
URL DataSource Disk Cache Size	The maximum disk cache size for binary content (attachments) sent as a part of the notification event.	Integer value. Default is 64.

PASW BIRT Report Designer

PASW BIRT Report Designer configuration options allow you to specify settings affecting the processing and display of reports. To modify the settings, click the corresponding option under PASW BIRT Report Designer in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-6
PASW BIRT Report Designer configuration options

Name	Description	Settings
BIRT Linked Resource Location	The directory on the server file system in which external resources for reports, like cascading stylesheets and images, are stored.	Full path of the directory containing external resources. To return to the system-defined default, click Use Default. This option restores the default directory that was established when you installed the system.
Enable SVG Chart	Specifies whether SVG chart output should be enabled. This setting should only be selected when SVG output is desired and the browsers viewing the report output are SVG enabled. If unselected, reports that would generate charts use the PNG image format instead of SVG.	Unselected by default.

Pager

The Pager Timeout configuration option allows you to specify the amount of time in minutes for paged data will be available. Changing this value may affect performance of the paging system. You must restart the repository for the new option value to take effect.

To modify the pager timeout:

- ▶ In the Configuration list, under Pager click Pager Timeout. The current value appears.
- ▶ In the Pager Timeout text box, enter the desired number of minutes.
- ▶ Click Set. The value you specified becomes the timeout value.
- ▶ To return to the system-defined default, click Use Default. This option restores the default value that was established when you installed the system.

Process Management

Process Management configuration options allow you to specify job execution settings as well as define the web service endpoints.

To modify the settings, click the corresponding option under Process Management in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-7
Process management configuration options

Name	Description	Settings
Calendar Pool	Duration that the process management server waits before repeating its scan of the repository for calendar schedules. Calendar schedules run based on their schedule time/date.	Integer value designating length of time in seconds. Default is 60.
Concurrency Idle	The number of seconds a thread may be idle before removing it from the pool. Threads are not removed from the pool if the pool size will fall below Concurrency Min.	Integer value designating length of time in seconds. Default is 600.
Concurrency Limit	The maximum number of threads that can be active at one time. Keep in mind that each actively running work (DBQ, for example) will consume one thread.	Integer value. Default is 25.
Concurrency Min	The minimum number of threads to create in the thread pool.	Integer value. Default is 5.
JMS Connection Factory Name	The name of JMS Connection Factory Name as registered with the JNDI service. Consult your Application Server documentation, or JMS server documentation for the appropriate value.	Default is <code>ConnectionFactory</code> . The name must be unique within the associated messaging provider.
JMS Naming Factory	The JMS Java class. For example, for JBoss application server the naming factory is <code>org.jnp.interfaces.NamingContextFactory</code> . This option can be set if all messages for all Message-based jobs are coming from a single remote server.	The default value is local application server JMS naming factory class name.
JMS Naming Service	The URI location of the naming service. For example, for JBoss application server the naming factory is <code>jnp://localhost:1099</code> . This option can be set if all messages for all Message-based jobs are coming from a single remote server.	The default value is local application server JMS naming service URI.
JMS Process Event Connection Factory	JMS connection factory class name to use for the process event queue.	The default value is local application server JMS naming factory class name.
JMS Process Event Queue	JNDI name of the JMS process event queue.	The default value is local application server JMS process event queue.
Maximum Number of Iterations	Maximum number of iterations for iterative job steps.	Integer value. Default is 250.

Name	Description	Settings
Message Poll	The length of time (in seconds) the Process Management server waits before repeating a scan of the Content Repository for message-based schedules that should be activated.	Integer value. Default is 120.
Process Notification Enabled	Indicates whether the Process Management server should interface with the Notification Server.	True or false. The default is true.
Submitted Work Expiration Time	The expiration period (in days) for submitted jobs.	Integer value. Default is 5.
Submitted Work Timestamp	Timestamp format to be used for Submitted Work folders.	Year, month, day, hour, minute, second format: yyyy.MM.dd.hh.mm.ss.SSS.
The date and time format for the time-stamped folders.	The date and time format for the time-stamped folders.	Year, month, day, hour, minute, second format: yyyy.MM.dd.hh.mm.ss.SSS.
The date format for the time-stamped folders.	The date format for the time-stamped folders.	Month, day, and year: MM-dd-yyyy.
The time format for the time-stamped folders.	The time format for the time-stamped folders.	Hour, minute, and second format: HH.mm.ss.

Reporting

The Reporting configuration option allows you to specify the path for writing out debugging information (as XML output) for visualization processing.

Important! If no value is specified for this option, debugging information for visualization processing is not generated.

To modify the directory path:

- ▶ In the Configuration list, under Reporting click Complete Visualization Directory. The current directory appears in the Complete Visualization Directory text box.
- ▶ Enter the new value of the absolute path of the directory.
- ▶ Click Set. The path that you specified becomes the default directory for writing out visualization processing information.

Repository

Repository configuration options allow you to define the Web service endpoints and toggle connection validation. To modify the settings, click the corresponding option under Repository in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-8
Repository configuration options

Name	Description	Settings
Categorical Value Limit	Limits the number of categorical variable values that are saved as PASW Modeler stream metadata. The limit is necessary to decrease the time it takes to save a stream to the repository.	Integer value. Default is -1 (no limit). Enter 0 to disable saving of values. Enter 1 or greater to limit the number of values saved.
Content Repository Endpoint	Defines Web service endpoint address for the content repository.	URL.
Credential passwords must be encrypted	Credentials passwords must be encrypted. False indicates that passwords can be passed as unencrypted text. <i>Note:</i> This option is redundant for PASW Collaboration and Deployment Services deployments where SSL is already enabled and should be used only in non-SSL deployments to encrypt credentials passwords.	Unselected by default.
Default Character Set	Defines the default character for the content downloaded from/uploaded to the server file system or when viewing repository files in a Web browser. The value is used only when the content such as a plain text file has not been explicitly assigned a character set.	A value designating the character set, such as UTF-8 or ASCII.
Log performance data	True indicates that performance data will be logged.	Unselected by default.
Message Bus Notification Enabled	Indicates whether the repository server should interface with the message bus.	Selected by default.
PASW Modeler Parameter Password Indicator	PASW Modeler stream parameters containing this string will be encrypted when stored and masked in the UI when a stream is scheduled for execution.	Masked password.
Reindex Queue Size	Defines the size of the queue to use for repository reindexing. This number should be greater than the value define by Reindex Thread Pool Size configuration option.	Integer value. Default is 15.
Reindex Thread Pool Size	Defines the number of threads to use for repository reindexing.	Integer value. Default is 5.
Repository Notification Enabled	Indicates whether the content repository server should interface with the notification server.	Unselected by default.

Name	Description	Settings
Resource Locking	Enables resource locking. Resource locking prevents a resource from being changed by multiple users at the same time. When enabled, a lock can be placed on a resource making the resource appear read only to others.	Selected by default.
Resource Transfer Lookup Table	Mapping implementation for ID lookup during resource transfers.	DISK or MEMORY.
Resource Transfer Page Result Cache Size	Size of the cache for storing page results during resource transfers. When the user performs individual conflict resolutions during resource transfer, there may be more conflicts than can be displayed at once in the user interface. The results cache size determines the number of pages cached for a single session. If the user is making heavy use of individual conflict resolution, it may help performance to increase the size of the cache; however, increasing the size of the cache will result in additional memory consumption.	Integer value. Default is 5.
Validate Server Executables	Specifies whether or not server executables should be validated when stored in the content repository.	Selected by default.

Scoring Service

The Scoring Service configuration options allow you to specify setting for the Scoring Service. To modify the settings, click the corresponding option under Scoring Service in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-9
Scoring Service configuration options

Name	Description	Settings
Audit Timer Period	The number of milliseconds between audit updates.	Integer value. Default is 3600000.
Default Logging destination.	Default logging destination.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the JMS queue for scoring logging.
Metrics Timer Period	The number of milliseconds between metric updates.	Integer value. Default is 5000.
Resolve Hostnames	Defines whether scoring service should attempt to resolve host names.	Selected by default.
Worker Pool Maximum Size	Maximum worker pool size.	Integer value. Default is 100.

Search

The Search configuration option allows you to specify the number of hits to display per page in Deployment Manager search results, result set size, as well as whether searches get logged in audit views. To modify the settings, click the corresponding option under Search in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-10
Search configuration options

Name	Description	Settings
Audit Searches	Log each search in the audit view. For more information, see the topic Auditing the Repository in Chapter 12 on p. 97. Note that enabling this option can slow down searches.	Unselected by default.
Default Page Size	Number of search results to display on a page.	Integer value. Default is 25.
Maximum Rows	Maximum number of rows in a search result set. The value must be set to -1 for unlimited number of results, or to a positive integer (to limit the size of the returned result set and avoid out of memory conditions or client timeout issues).	Integer value. Default is -1.

Security

Security configuration options allow you to specify repository access settings.

To modify the settings, click the corresponding option under Security in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-11
Security configuration options

Name	Description	Settings
Cache Logins	Saves logins for faster response from Web services. If enabled, changes to users, groups or roles will take 30 minutes or longer to become effective. Requires a server restart.	Selected by default.
Cache Session Timeout	Number of minutes before an idle user's security session is removed.	Integer value. Default is 30.
Disable Clients	Disables login for PASW Collaboration and Deployment Services client applications (Deployment Manager, Deployment Portal, etc.)	Unselected by default.

Name	Description	Settings
Encrypt Password	Requires Web services to use encrypted passwords. Web services will send an encryption key when requesting passwords. The server will encrypt the password using the public key provided. If Encrypt Password is selected, Web services will not be allowed to request passwords by providing an encryption key. This affects user preferences, content repository credentials, and similar services.	Selected by default.
Lowercase User ID	Forces the internal identifier for a user to be lowercase. This should be unselected only if a remote user directory depends on case-sensitive user IDs.	Selected by default.
Message	Message appearing on the Deployment Manager welcome screen.	Message text. HTML tags can be used to apply formatting.
Normalize Principal	Specifies that user names are stored in the database in normalized character format when users are created or imported (<i>Normalization Form C</i> as defined by the Unicode technical standard)	Unselected by default.

Setup

The Setup configuration option allows you to specify miscellaneous setup setting for the repository, such as the URL prefix used in references to PASW Collaboration and Deployment Services, JMS queue setting, and JMS message bus settings.

To modify the settings, click the corresponding option under Setup in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-12
Security configuration options

Name	Description	Settings
Log JMS Connection Factory	JNDI name of the log JMS connection factory.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the log JMS connection factory.
Log JMS Queue	JNDI name of the log JMS queue.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the log JMS queue.

Name	Description	Settings
Message Bus JMS Connection Factory	JNDI name of the message bus JMS connection factory.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the message bus JMS connection factory.
Message Bus JMS Topic	JNDI name of the message bus JMS topic.	A deployment-specific and/or server-specific case-sensitive string used by the JNDI service to identify the message bus JMS topic.
URL Prefix	The prefix should be resolvable in DNS (or WINS). If using SSL, the prefix should begin with <i>https</i> instead of <i>http</i> . Furthermore, the port can be omitted if the server uses the standard <i>http</i> port of 80, or the standard <i>https</i> port of 443. The server must be restarted for any changes to the prefix to take effect.	URL.

ShowCase

ShowCase configuration options allow you to specify connection settings used for creating ShowCase Warehouse Builder job steps in Deployment Manager (ShowCase Warehouse Builder is separately licensed and installed with the ShowCase Suite product set).

To modify the settings, click the corresponding option under Showcase in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 7-13
ShowCase configuration options

Name	Description	Settings
ShowCase Warehouse Builder Database	The database/library name for ShowCase Warehouse Builder set and definition information.	Valid database/library name.
ShowCase Warehouse Builder Host Name	The IP address or hostname of the IBM i server used by Warehouse Builder.	Any valid IP address or hostname.
ShowCase Warehouse Builder User name	The username for connecting to the database/library mentioned above. This is applicable only when the PASW Collaboration and Deployment Services server is installed on Windows.	Username.
ShowCase Warehouse Builder User Password	The password for connecting to the database/library mentioned above. This is applicable only when the PASW Collaboration and Deployment Services server is installed on Windows.	Masked password.

MIME Types

Multipurpose Internet Mail Extensions, or **MIME**, is a standard for identifying different types of information. MIME originated as an extension of e-mail, but it is also used by HTTP to define the content being delivered by a server.

When responding to a request for a file, a server appends header information to the file. This information includes the MIME type, denoting the media type contained within the file. The server uses the extension of the file to define the MIME type. The client receiving the file uses the MIME type to determine the best method for handling the file.

The server controls the associations between file extensions and MIME types. To configure these mappings, use the MIME Types and File Type Icons page of Deployment Manager, accessed by clicking MIME Types in the navigation list.

Figure 8-1
MIME Types and File Type Icons page

Name	MIME Type	Extensions	Small Icon	Delete
Application Server Data Source	application/x-vnd.spss-datasource-appserver			
Application View	application/x-vnd.spss-application-view	av		
Bitmap Image	image/bmp	bmp		
Cascading Style Sheet	text/css	css		
Complex Samples Analysis Plan	application/x-vnd.spss-spss-csaplan	csaplan		
Complex Samples Plan	application/x-vnd.spss-spss-csplan	csplan		
Custom Dialog Package	application/x-vnd.spss-statistics-spd	spd		
Data Provider Definition	application/x-vnd.spss-data-provider	dpd		
Data Service Data Source	application/x-vnd.spss-realtime-dataservice			
Domain	application/x-vnd.spss-repository-domain	CredentialRealm		
Enterprise View	application/x-vnd.spss-enterprise-view			
Extended (Enhanced) Windows Metafile Format	image/x-ernf	ernf		
Extensible Markup Language File	text/xml	xml		
Folder	application/x-vnd.spss-repository-folder	Folder		
Graphics Interchange Format Image	image/gif	gif		

On the MIME Types and File Type Icons page, you can perform the following tasks:

- Add MIME type mappings to the server.
- Edit existing MIME type settings, including the assignment of images to files.
- Delete MIME type mappings from the server.

Note: Many common icons do not appear in Deployment Portal by default. For external file types (for example, *application/msword*), administrators can assign an icon to the MIME type. For more information, see the topic [Adding MIME Type Mappings](#) on p. 73.

Adding MIME Type Mappings

A MIME type consists of two parts, a type and a subtype, separated by a forward slash. The type specifies the general media type as *application*, *audio*, *image*, *message*, *model*, *multipart*, *text*, or *video*. The subtype, on the other hand, identifies the format for the media and varies across media types. For example, *text/html* corresponds to text in HTML format.

Subtypes often include prefixes to identify MIME types for specific products. For example, subtypes associated with commercial products include the prefix *vnd.*, designating a vendor subtype, such as *application/vnd.ms-access*. In contrast, subtypes for noncommercial products include the prefix *prs.*, denoting a personal subtype.

MIME types should be registered with the Internet Assigned Numbers Authority (IANA). Types that are not registered should prefix the subtype with *x-* to prevent conflicts with types that may be registered in the future, as in *application/x-vnd.spss-clementine-stream*. For a list of registered MIME types, consult the [IANA](http://www.iana.org/assignments/media-types/) (<http://www.iana.org/assignments/media-types/>).

To add a new MIME type mapping:

- ▶ On the MIME Types and File Type Icons page, click Add New MIME Type. The Add MIME Types and File Type Icons page appears.

Figure 8-2
Creating MIME types

The screenshot shows a web-based dialog box titled "Add MIME Types and File Type Icons". It contains the following fields and controls:

- Name:** A text input field.
- MIME Type:** A text input field.
- Extensions:** A text input field with a note below it: "To enter multiple extensions, separate extensions with spaces."
- Small Icon:** A radio button selected for "No", and a "Browse..." button.
- Buttons:** "Save" and "Cancel" buttons at the bottom right.

- ▶ Enter a name for the MIME type. The name provides an identifier of the type that is easier to read than the type itself. For example, the name *Custom Dialog Package* is easier to read than the type *application/x-vnd.spss-statistics-spd*.
- ▶ Enter the MIME type being added.
- ▶ Enter the file extensions to associate with the MIME type. Use a space between entries when specifying multiple extensions.

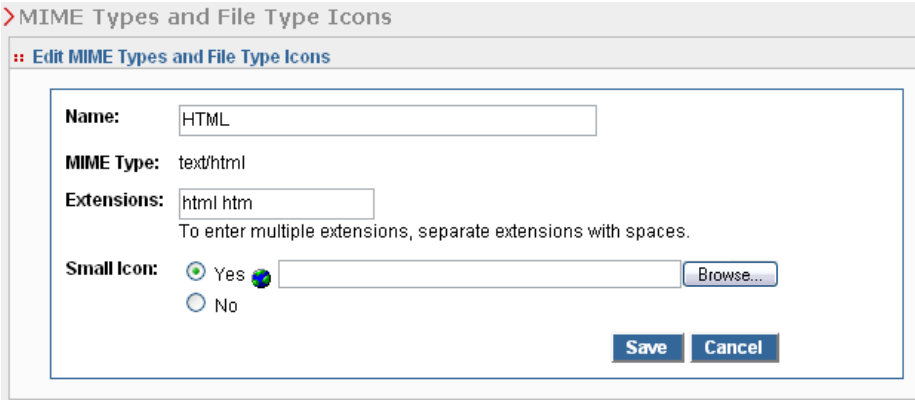
- ▶ Assign an icon to the MIME type. This image should be 16 x 16 pixels in size and must be a *.gif* file. The image is typically used in content lists. Click Browse to navigate to the file. If no icon assignment is desired, select No.
- ▶ Click Save to add the MIME type and return to the Add MIME Types and File Type Icons page, or click Cancel to return without saving the MIME type to the server.

Editing MIME Type Mappings

To edit an existing MIME type:

- ▶ On the MIME Types and File Type Icons page, click the name of the MIME type to be edited. The Edit MIME Types and File Type Icons page for that MIME type appears.

Figure 8-3
Editing MIME types



The screenshot shows a web browser window with the address bar displaying '>MIME Types and File Type Icons'. The main content area shows a dialog box titled 'Edit MIME Types and File Type Icons'. Inside the dialog, there are several input fields and controls:

- Name:** A text box containing 'HTML'.
- MIME Type:** A text box containing 'text/html'.
- Extensions:** A text box containing 'html htm'. Below it, a note says 'To enter multiple extensions, separate extensions with spaces.'
- Small Icon:** A section with two radio buttons: 'Yes' (which is selected) and 'No'. To the right of the 'Yes' radio button is a small globe icon and a text box. To the right of the text box is a 'Browse...' button.
- At the bottom right of the dialog are two buttons: 'Save' and 'Cancel'.

- ▶ Modify the settings as desired. Icons will be changed only if you select a new file or select No. To delete an icon, select No.
- ▶ Click Save to save the new settings for the MIME type and return to the Add MIME Types and File Type Icons page, or click Cancel to return without saving the new MIME type settings to the server.

Deleting MIME Type Mappings

To delete an existing MIME type:

- ▶ On the MIME Types and File Type Icons page, click the delete icon for the MIME type to be deleted.

The MIME type table refreshes, reappearing without the deleted MIME type.

Reindexing the Repository

Indexing is used to optimize repository search. By default, when the repository is upgraded, the old index is cleared and the index is rebuilt. The repository can also be configured to force reindexing of processing results, such as job output, at startup. For more information, see the topic [Process Management](#) in Chapter 7 on p. 64. The repository search is automatically disabled while reindexing is run at startup.


Reindexing can also be performed on demand in the browser-based Deployment Manager by an authorized user. For more information, see the topic [Actions](#) in Chapter 4 on p. 36.

Note: Reindexing is a resource-intensive and lengthy process that should be run only when it is absolutely necessary, such as when a lot of new data are imported into the repository. It is strongly recommended that reindexing be run only when there is no user activity in PASW Collaboration and Deployment Services. If it impossible to ensure that all users are logged out at the time reindexing is run, repository search must be disabled; however, it is not advised to clear the index if the system is being used.

To reindex repository:

1. In the browser-based Deployment Manager, click Repository Index in the navigation list. The Content Repository Indexing page appears.

Figure 9-1
Content repository indexing



2. Do one of the following:
 - If no users are logged in to repository, select Clear the entire index before reindexing.
 - If users are still logged in to repository, select Disable Clients while indexing is running.
3. Click Start Indexing. While the index is being rebuilt, the Content Repository Indexing Status page displays the statistics of processed objects.

Figure 9-2
Content repository indexing status

> Content Repository	
:: Indexing Status	
Status:	Complete
Search Disabled:	Yes
Clear Index:	Yes
Start Time:	Oct 20 2006 11:35:23 AM
End Time:	Oct 20 2006 11:35:23 AM
Elapsed Time:	00:00:00 (hh:mm:ss)
Folders Indexed:	2
Topics Indexed:	1
Files Indexed:	0
Velocity:	∞ objects/second

Notifications

PASW Collaboration and Deployment Services provides the mechanisms of **notifications** and **subscriptions** for keeping the users informed about changes to repository objects and job processing results. Both notifications and subscriptions generate e-mail messages when corresponding events occur. For example, when a job fails, PASW Collaboration and Deployment Services can automatically send an e-mail to the person responsible for the job. The failure triggers a search for a template matching the event. Applying the template to the event creates an e-mail that is sent to any recipients associated with the event.

Notification templates included in the default repository installation are located in the subdirectories of *<Installation Directory>\components\notification\templates*. The names of the subdirectories correspond to the general event type. For example, the folder *components\notification\templates\PRMS\Completion* contains two message templates. These templates, *job_success.xml* and *job_failure.xml*, correspond to the success and failure, respectively, of job executions. If a job completes successfully, PASW Collaboration and Deployment Services uses the *job_success* template to generate a notification message indicating that success. The content and appearance of the notification messages can be customized by modifying the templates.

Notification Message Template Structure

Notification templates transform event information into notification messages using Apache **Velocity** Template Language.

Velocity Template Structure

A Velocity template has a *.vm file extension. The template generates a message using the = operator to assign the /mimeMessage/messageSubject, /mimeMessage/messageContent, and /mimeMessage/messageProperty values that are subsequently parsed by the e-mail processor. The following sample template generates a simple, generic e-mail message indicating the success of the corresponding job.

```
/mimeMessage/messageSubject=Job Completion  
/mimeMessage/messageContent[text/plain; charset=utf-8]=The job completed successfully.
```

The following figure displays the resulting e-mail.

Figure 10-1
Generic notification message

The job completed successfully.

For more information about Velocity templates, see the Apache [Velocity project](http://velocity.apache.org/) (<http://velocity.apache.org/>) documentation.

Message Properties

E-mail notification templates may include properties that determine how a message is processed in cases where SMTP settings different from repository defaults are to be used. For example, it may be necessary to specify a different SMTP server name and port number or the return e-mail address assigned to the message. Default SMTP properties are listed under repository notification configuration options. For more information, see the topic [Notification](#) in Chapter 7 on p. 60. If Sun JVM is used with the repository installation, SMTP properties will correspond to the JavaMail API properties for message handling defined in the “[Message properties](#)

” table. Note that these properties may be different for different Java environments. For detailed information about SMTP properties, see the JVM vendor documentation.

Table 10-1
Message properties

Message Property	Attribute	Event Property	Description
mail.debug	value	MailSmtpDebug	A Boolean value indicating the initial debug mode. The default is false.
mail.smtp.user	value	MailSmtpUser	The default SMTP username.
mail.smtp.password	value	MailSmtpPassword	The SMTP user password.
mail.smtp.host	value	MailSmtpHost	The SMTP server to which to connect.
mail.smtp.port	value	MailSmtpPort	The SMTP server port to which to connect. The default is 25.
mail.smtp.connectiontimeout	value	MailSmtpConnectionTimeout	The socket connection timeout value in milliseconds. By default, the timeout is infinite.
	value	MailSmtpTimeout	The socket I/O timeout value in milliseconds. By default, the timeout is infinite.
mail.smtp.from	value	MailSmtpFrom	The e-mail address used for the SMTP MAIL command. This sets the envelope return address.
mail.smtp.from	label	MailSmtpFromPersonal	The envelope return address label.
mail.smtp.localhost	value	MailSmtpLocalhost	The local hostname. The property should not normally need to be assigned if the JDK and name service are configured properly.

Message Property	Attribute	Event Property	Description
mail.smtp.ehlo	value	MailSmtpEhlo	A Boolean value indicating whether or not to sign on with the EHLO command. The default is true. Typically, failure of the EHLO command results in a fallback to the HELO command. This property should be used only for servers that do not fall back.
mail.smtp.auth	value	MailSmtpAuth	A Boolean value indicating whether or not to authenticate the user using the AUTH command. The default is false.
mail.smtp.dsn.notify	value	MailSmtpDsnNotify	Specifies the conditions under which the SMTP server should send delivery status notifications to the message sender. Valid values include: <ul style="list-style-type: none"> ■ NEVER indicates that no notification should be sent. ■ SUCCESS indicates that a notification should be sent on successful delivery only. ■ FAILURE indicates that a notification should be sent on a failed delivery only. ■ DELAY indicates that a notification should be sent when the message is delayed. Multiple values can be specified using a comma separator.

The syntax for defining these properties in a Velocity template is as follows:

- The property value must be assigned to `mimeMessage/messageProperty` with property name and label arguments in square brackets, as in the following example:

```
/mimeMessage/messageProperty[smtp.mail.smtp.from][Brian McGee]=bmagee@mycompany.com
```

- The value of property label is optional; therefore, the assignment statement can have the following syntax:

```
/mimeMessage/messageProperty[smtp.mail.smtp.from][]=bmagee@mycompany.com
```

- The values of property name and label can be assigned as static values or through variables referencing the corresponding event properties:

```
/mimeMessage/messageProperty[smtp.mail.smtp.from][$MailSmtpFromPersonal]=$MailSmtpFrom
```

Message Content

The content of a notification message corresponds to the text supplied for the `messageSubject` and `messageContent` elements of the notification template. For either element, this text may include variable event property values.

- In Velocity templates, variable values are referenced using the `$` notation. For example, `Job step ${JobName}/${JobStepName} failed at ${JobStepEnd}` inserts the text with the current values for the `JobName`, `JobStepName`, and `JobStepEnd` properties.

The variables that can be inserted into a message reference the properties of the event that triggers the notification. Typical properties include:

- `JobName`, a string denoting the name of the job.
- `JobStart`, a timestamp indicating the time the job began.
- `JobEnd`, a timestamp indicating the time the job ended.
- `JobSuccess`, a Boolean value indicating whether or not the job was successful.
- `JobStatusURL`, a string corresponding to the URL at which the job status can be found.
- `JobStepName`, a string denoting the name of the job.
- `JobStepEnd`, a timestamp indicating the time the job ended.
- `JobStepArtifacts`, an array of string values denoting the URLs of the job step output.
- `JobStepStatusURL`, a string corresponding to the URL at which the job step status can be found.
- `ResourceName`, a string corresponding to the name of the object affected by the event, such as the file or folder name.
- `ResourcePath`, a string corresponding to the path of the object affected by the event.
- `ResourceHttpUrl`, a string corresponding to the HTTP URL at which the object can be found.
- `ChildName`, a string corresponding to the name of the child object of the parent object affected by the event. For example, when a file is created in a folder, this will be the name of the file.
- `ChildHttpUrl`, a string corresponding to the HTTP URL at which the child object can be found.
- `ActionType`, for repository events, the type of action that generated the event—for example, `FolderCreated`.

The available properties are defined by the event and will be different for different event types.

The following sample Velocity template for job step success notification inserts the names of the job and job step in the subject line. The content of the message also includes the end times for the step, the URL at which the status can be viewed, and a list of artifacts generated by the job step. Note that the template uses the `#foreach` loop structure to retrieve the URLs of the artifacts from the `JobStepArtifacts` property array.

```
/mimeMessage/messageSubject=Job step ${JobName}/${JobStepName} completed successfully
/mimeMessage/messageContent[text/html;charset=utf-8]=
<html>
<head>
<meta http-equiv="Content-Type" content="text/html;charset=utf-8"/>
```



```

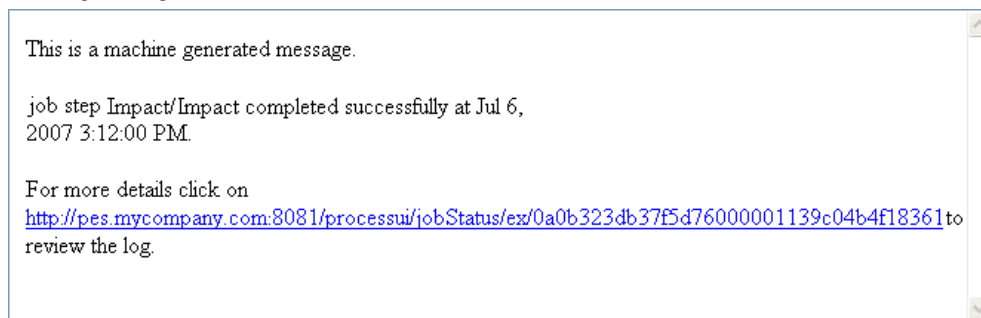
</head>
<body>
<p>This is a machine generated message.</p>
<p>job step ${JobName}/${JobStepName}
completed successfully at ${JobStepEnd}.</p>
<p>For more details click on <a href="${JobStepStatusURL}">${JobStepStatusURL}</a>
to review the log.</p>
#if ($JobStepArtifacts)
#foreach($artifact in $JobStepArtifacts)
  <a href='${artifact.get("url")}'>${artifact.get("filename")}&lt;</a><br>
#end </body>
#end
</html>

```

The following figure displays the resulting e-mail.

Figure 10-2

Message using customized content



The following code segments demonstrate how the Velocity template for folder content notification can be modified to remove the hyperlink to the job from the message. PASW Collaboration and Deployment Services jobs cannot be opened outside Deployment Manager; therefore, it is strongly recommended to customize the notification message to remove the hyperlink. The additional if-condition in the example tests the MIME type of the object; if the object is an PASW Collaboration and Deployment Services job, the hyperlink is not included.

Original template:

```

#if($Attachments)
See attachment.
#else
<p>To review the content of the file, go to <a href='${ResourceHttpUrl}'>${ResourceHttpUrl}</a>.</p>
#end

```

Modified template:

```

#if($Attachments)
See attachment.
#else
#if($MimeType!='application/x-vnd.spss-prms-job')
<p>To review the content of the file, go to <a href='${ResourceHttpUrl}'>${ResourceHttpUrl}</a>.</p>
#end

```

```
#end
```

Message Format

A notification template must specify the MIME type of the message content. In notification templates, the MIME type argument is specified in square brackets with `/mimeMessage/messageContent`.

The MIME type can have one of two values:

- *text/plain*. Notification messages appear in plain text. This is the default setting.
- *text/html*. Notification messages include HTML tags. Use this setting to control the appearance of the content within the message. The HTML within the message must be well-formed.

It is a good practice to always encode template output as Unicode (UTF-8).

HTML notification templates can take advantage of the functionality allowed in the markup. For example, the message can include a link to a Web page or to output from the job.

The following template generates a notification message for job step completion, formats content as a table, specifies background color for the message using an inline style for body, and defines a blue Verdana font for paragraphs using an internal style sheet. The message also includes a link to the job output.

```
/mimeMessage/messageSubject=${JobName}/${JobStepName} completed successfully
/mimeMessage/messageContent[text/html;charset=utf-8]=
  <html>
  <head>
  <meta http-equiv="Content-Type" content="text/html;charset=utf-8"/>
  <style type="text/css">
  table {font-family: verdana; color: #000080}
  p {font-family: verdana; color: #000080}
  .foot {font-size: 75%; font-style: italic} </style>
  </head>
  <body style="background-color: #DCDCDC">
  <table border="8" align="center" width = 100%>
  <tr align="left">
  <th>Job/step name</th>
  <td>${JobName}/${JobStepName}</td>
  </tr>
  <tr align="left">
  <th>End time</th>
  <td> ${JobStepEnd}</td>
  </tr>
  <tr align="left">
  <th>Output</th>
  <td><p>
  #if ($JobStepArtifacts)
  #foreach($artifact in $JobStepArtifacts)
  <a href='${artifact.get("url")}'>${artifact.get("filename")}</a><br>
  #end
  #else None <br>
  #end
```

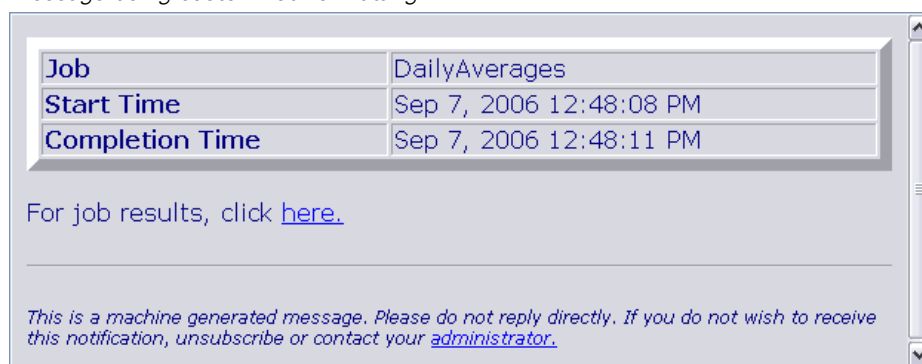
```

<p></td>
</tr>
</table>
<hr/>
<p class="foot">This is a machine generated message.
Please do not reply directly. If you do not wish to receive
this notification, unsubscribe or contact your
<a href="mailto:admin@mycompany.com"> your PASW Deployment
Services administrator.</a></p></body>
</html>

```

The following figure displays the resulting e-mail.

Figure 10-3
Message using customized formatting



Editing Notification Templates

To edit a Velocity message template:

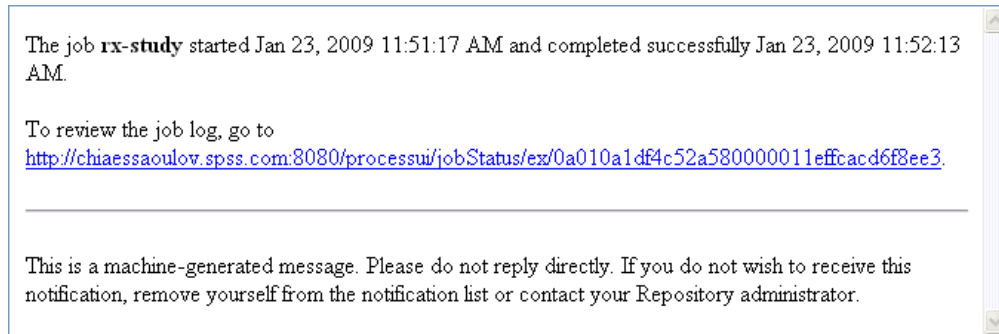
1. Open the template in a text editor. Subfolders of the *components/notification/templates* folder contain the current set of templates in use.
2. Modify the value assigned to `/mimeMessage/messageSubject`. Use the `$` notation to insert event property variables into the message subject. For more information, see the topic [Message Content](#) on p. 80.
3. Define the MIME type of the message. The MIME type value is specified in the square brackets following `messageContent`. For a plain text message, use a value of *text/plain*. For an HTML message, use a value of *text/html*. For more information, see the topic [Message Format](#) on p. 82.
4. Modify the value assigned to `messageContent`. Use the `$` notation to insert event property variables into the message content.
5. Save the template using its original name.

Subsequent notification messages will use the modified templates when the corresponding event occurs.

Job Status

A notification template that includes the *JobStatusURL* property yields a message containing a link to the job output and log.

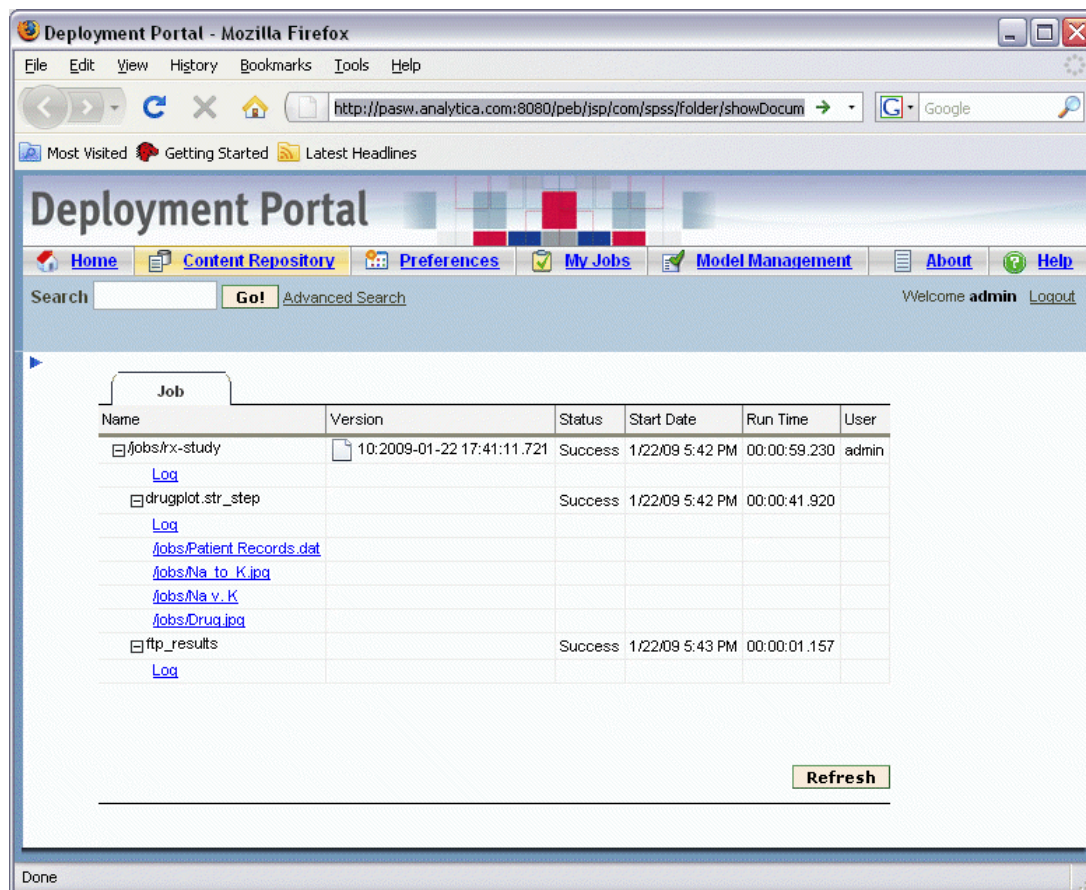
Figure 10-4
Message containing status link



To view the results of a job:

1. Click the status link in a notification message. The Login page for the server opens.
2. Enter your login name and password. Click Login. The Job Status page opens.

Figure 10-5
Job status



Jobs status view displays the processing status details of a job, including the information about the status of all job steps in the job. Using the view, you can display the job log, the logs of individual job steps, as well as the generated output.

Job Details

Name. The repository path of the job.

Version. The version label of the job.

Status. The processing status of the job, such as *Running*, *Succeeded*, or *Failed*.

Start Date. The date and time the job processing started.

Run Time. The duration of job execution.

User. The user who submitted the job.

- ▶ To refresh the status of the job, click Refresh.

- ▶ To expand the details for the job, which include job log and job steps, click + next to the job name.
- ▶ To display the job log, click Log link under the job name. The Log tab opens. To close the tab, click Close.

Job Step Details

Name. The name of the job step.

Status. The processing status of the job step, such as *Running*, *Succeeded*, or *Failed*.


Start Date. The date and time the job step processing started.

Run Time. The duration of job step execution.

- ▶ To expand the details for a job step, which include job step log and any resulting output, click + next to the job step name.
- ▶ To display the job step log, click Log link under the job step name. The job step log opens on a new tab. To close the tab, click Close.
- ▶ To display job step output, click the output file name. The Results tab opens. To close the tab, click Close.

For example, the job shown here consists of two steps. The first step involves a PASW Modeler stream, *drugplot.str*, which took 59 seconds to complete and generated four outputs. The second step called a script to transfer the generated files to a remote host, which took just one seconds.

Figure 10-6
Job output



Optimizing Notification Service Performance

The overall performance of the notification service is a combination of the performance of PASW Collaboration and Deployment Services components that manage subscriber and subscription data, collect events, and generate, format, and distribute notifications, as well as the performance of the database system that stores and processes the subscription data. Notification functions of PASW Collaboration and Deployment Services require significant system resources and may need to be fine-tuned. It is also recommended to follow the general guidelines for notification service performance improvement.

Notification Service Configuration

Notification configuration options

Notification service performance may be improved by changing the parameters defined by the notification service configuration options. The following options may have a noticeable positive effect on performance:

- Event noise filtering enables the system to ignore notification events that do not have matching subscriptions with subscribers or associated notification providers early in the process. Event noise filter cache size defines the maximum number of cached events that do not resolve in any matching subscriptions. Enabling event noise filtering (*Event Noise Filter* configuration option) and, if necessary, increasing the size of the cache (*Event Noise Filter Cache* configuration option) can improve notification service performance. Disabling event noise filtering is not recommended in the production environments and should be used only for debugging and testing purposes.
- Subscription identifiers cache is a cache of mappings for the resolved filtering expressions to the list of matching subscription identifiers. The size of the cache defines the number of the filtering expressions in the cache. While there is no limitation on the number of matching subscription identifiers associated with the filtering expressions, it is expected that the number of matching subscriptions per resolved filtering expression would be relatively small—for example, a few dozen or, in rare cases, several hundred. Increasing the size of the cache (*Subscription Identifiers Cache* configuration option) can improve performance.
- Persistent event queue enables the system to maintain a cache of incoming notification events in temporary disk storage to minimize the amount of consumed memory. By default, incoming notification events are kept in memory. If the rate of the incoming events is high and the amount of the available RAM is not sufficient, it is possible to store events in the temporary disk storage. If the persistent event queue is enabled, the event queue storage commit batch size sets the maximum number of notification events to be kept in memory before writing them out to temporary storage. While enabling the persistent event queue (*Persistent Event Queue Enabled* configuration option) and increasing the commit batch size (*Persistent Event Queue Size* configuration option) can improve performance, only moderate increases in batch size are recommended because of additional memory requirements. Increasing the size of the persistent event queue storage file on the disk (*Persistent Event Queue Size* option) does not visibly affect performance. Note that the system must be restarted for the changes to the persistent event queue settings to take effect.
- Disabling binary content (e-mail attachments) sent with the notification message can significantly improve performance (*Binary Content Enabled* configuration option). Generating of the notification messages with binary attachments can be a processing-intensive operation. The content of the binary attachment must be read from the repository, added to the notification message, and routed through the appropriate distribution channel, such as an e-mail server. Some transformation of the binary content of the attachment may also be required for particular types of notification messages. For example, base-64 encoded binary attachments (SMTP) will add about 33% to the total size of the generated messages. Processing overhead can be even greater if a number of different custom templates are used to format notification messages with large attachments. In these cases, the notification service must format messages, add attachments, and push each message through the distribution channel separately. In order to improve performance, it is advisable to limit the number

of notifications with attachments, the size of the attachments, and the number of custom templates that will be used to format notification messages with attachments.

- The processing and distribution of notification messages is very resource-intensive. For smaller installations, or when PASW Collaboration and Deployment Services is installed on a non-dedicated server, it is advisable to limit the size of the pool to a single background thread by modifying the *Core Event Collector Pool Size* and *Maximum Event Collector Pool Size* configuration options.

For a complete listing of notification configuration options, detailed descriptions, and default values, see [Notification on p. 60](#)

Dedicated SMTP Server

The performance of the delivery channel, such as an e-mail server, is the critical factor controlling the overall performance of the notification service. For PASW Collaboration and Deployment Services notifications, it is strongly recommended to use a fast, dedicated SMTP server rather than the regular corporate e-mail server. Using a dedicated server has been demonstrated to dramatically reduce the time it takes to add a notification message to the mailer queue, thus significantly improving the performance of the notification service. One possible configuration is deploying a dedicated e-mail server on the same host as the repository, which reduces the time it takes notification service to communicate with the e-mail server over the network.

Number of Threads

It is essential that the number of threads allocated by the SMTP server is sufficient. The number must be equal to or greater than the number of processing threads in the event collector pool of the PASW Collaboration and Deployment Services notification service. If the distribution server has an insufficient number of threads, the notification service will not be able to communicate with it efficiently.

General Recommendations

Using the following techniques can significantly improve the performance of the notification service without reducing the overall functionality available to the PASW Collaboration and Deployment Services user.

Minimize the number of recipients.

To minimize the overall recipient aggregation time during event matching, it is advisable to define a set of external distribution lists instead of specifying each subscriber individually. These distribution lists can be maintained in corporate directory servers (Microsoft Exchange, Lotus Domino, etc.). This approach eliminates the need for the rather large number of database queries that the notification service must perform to retrieve recipients and their delivery devices. Specialized corporate SMTP servers should be able to utilize available resources and handle delivery of the notification messages much more efficiently.

Minimize the number of custom templates.

PASW Collaboration and Deployment Services provides the capability to define an unlimited number of custom templates that will be used to format notification messages for a given event type. However, under normal circumstances, it is sufficient to format notification messages using only the default templates. The default templates are stored in the file system on the server and cached in memory. These templates can be customized to meet specific user requirements. For more information, see the topic [Editing Notification Templates](#) on p. 83. A large number of custom templates (hundreds or thousands per matching event) can visibly degrade performance because the templates must be retrieved from the database on each request and each notification message should be formatted separately. The same rationale applies to a custom SMTP From address. In most cases, it is sufficient to have a single default From address specified as a repository configuration option. Even if the content (subject and body) of the notification template is the same as that of the default template, specifying a custom From address establishes a custom template for a given notification.

Minimize the number of subscriptions.

To improve performance of the notification service, it is generally desirable to minimize the number of subscriptions that will be matched by a single event. If the incoming event matches a large number of subscriptions that have different subscribers and different message templates, the system will not be able to efficiently aggregate the distribution and will have to generate separate notification messages to the recipients. It is important to note that a single initial notification event can generate a number of derived events as processing traverses the event type hierarchy. An initial event can also be broken out into a series of events by application-specific event splitters. If a large number of derived events will be generated for an initial event, it is advisable to come up with a strategy for managing subscription layouts. For example, instead of specifying a number of separate subscriptions for each child folder in the content repository hierarchy, it is often sufficient to specify a single subscription for the parent folder and use the Apply to Subfolders option. For more information, see the Deployment Manager user documentation. Limiting the number of individual subscriptions can be also beneficial. Instead of allowing users to subscribe individually, distribution lists can be set up and maintained on corporate SMTP servers. Distribution lists can be used to create a limited number of subscriptions in order to improve performance and minimize message processing and distribution time.

Schedule subscription management activities.

To improve performance during event matching, the PASW Collaboration and Deployment Services notification service maintains a number of internal caches. These caches are invalidated (cleared) if the client makes modifications to the event type repository or the subscription repository. It is advisable to perform subscription management activities, such as adding subscribers, deleting subscriptions, etc., based on a schedule that does not overlap with the peak event processing times for the notification service. Performing subscription management activities under a light processing load is generally acceptable but can lead to short bursts of poor performance.

Debugging the Notification Service

To enable debugging for the notification service, edit the *log4j.xml* file of your application server. If you are using JBOSS, enable DEBUG logging level for the *com.spss.notification* package by editing `<your_jboss_installation>\server\default\conf\log4j.xml` as follows:

```
<category name="com.spss.notification"> <priority value="DEBUG"/> </category>
```

Other application servers can provide browser interfaces or some other ways of editing logging configuration for the deployed components. To enable SMTP logging, set the *SMTP Turn on Debug Mode* configuration option to true in Deployment Manager. While the notification log is very verbose and provides very detailed information about event matching and notification distribution activities, the most important log item to look for is:

```
[...SmtpDistributor] Exiting SMTP distributor. The distribution took 5.906 s.
```

If the SMTP distribution takes more than 100–200 milliseconds, it is strongly recommended to use a dedicated SMTP server.

For debugging purposes, it is also advisable to enable Delivery Status Notifications (DSN) by setting the corresponding configuration option to the following values:

SMTP DSN Notify

FAILURE,SUCCESS,DELAY

SMTP DSN Ret

FULL

Note: Your SMTP server must support the RFC3461 specification to generate these delivery notifications.

Troubleshooting Notification Delivery Failures

If correct settings have been specified for the e-mail server and the default sender's e-mail address during the installation of repository, additional e-mail configuration is not usually required in order for PASW Collaboration and Deployment Services notifications to be delivered successfully. If a mistake has been made during the installation, it can be corrected by changing notification configuration options. For more information, see the topic [Notification](#) in Chapter 7 on p. 60.

PASW Collaboration and Deployment Services administrator is also notified when delivery failures for notifications and subscriptions with a system-generated message similar to the following:

Your message did not reach some or all of the intended recipients.

Subject: PASW Deployment Services: New version of ChurnAnalysis created
Sent: 4/5/2007 2:35 PM

The following recipient(s) could not be reached:

jsmiht@mycompany.com on 4/5/2007 2:35 PM

There was a SMTP communication problem with the recipient's email server.
Please contact your system administrator.

In most cases, delivery failures are caused by user error when specifying notification recipients or default subscription addresses.

In certain cases, it is possible to experience problems with the delivery of notification messages due to the setup of the corporate network or the e-mail server. For example, the server may not be configured to relay to external addresses. The following steps can be taken to investigate the problem:

- To definitively diagnose notification delivery failures, use repository audit records. For information about auditing, see Chapter 12.
- To determine the cause of the notification failure, it is recommended to enable the debugging mode. For more information, see the topic [Debugging the Notification Service](#) on p. 91.
- *nslookup* queries can be used to examine the configuration of your SMTP server.
- Examining the SMTP headers of the notification messages can provide useful information about SMTP server message relaying.

Notification and subscription delivery failures are logged in repository auditing views. For more information, see the topic [Auditing the Repository](#) in Chapter 12 on p. 97.

JMS Setup

PASW Collaboration and Deployment Services uses Java Messaging Service (JMS) to communicate with third-party applications and trigger job processing based on repository events. The JMS API is a Java Message Oriented Middleware (MOM) API for sending messages between two or more clients. Using JMS, a program first creates an instance of a connection factory to connect to the queue or topic and then populates and sends or publishes the messages. On the receiving side, the clients then receive or subscribe to the messages. The same Java classes can be used to communicate with different JMS providers by using the JNDI information for the desired provider.

PASW Collaboration and Deployment Services supports JMS communication based only on the publish/subscribe model, in which messages are published to a particular message topic. Zero or more subscribers may register interest in receiving messages on a particular message topic. JMS queues are currently not supported.

The procedure for setting up JMS to work with PASW Collaboration and Deployment Services will be different depending on the JMS provider used by a specific PASW Collaboration and Deployment Services installation. Most popular open-source JMS providers include Apache ActiveMQ, OpenJMS from the OpenJMS Group, and JBoss Messaging from JBoss. Proprietary implementations include WebSphere MQ from IBM (formerly MQSeries), SAP NetWeaver WebAS Java JMS from SAP AG, Oracle AQ, Sun Java System Message Queue, and BEA WebLogic.

The examples below demonstrate the use of JMS in PASW Collaboration and Deployment Services: The first example illustrates how to configure a sample program to communicate with PASW Collaboration and Deployment Services using JBoss Messaging; the example provides a JMS topic configuration file, the source code of a Java class that will be publishing JMS messages to the topic, and a Windows batch file to run the class. The second example shows how to set up job processing based on repository events. For information about setting up other JMS providers, see the vendor documentation.

JMS Topic Configuration

In JBoss, JMS queues and topics are known as **administered objects**, and they can be deployed through configuration files—for example:

```
<?xml version="1.0" encoding="UTF-8"?>
<server>
  <mbean code="org.jboss.mq.server.jmx.Topic"
    name="jboss.mq.destination:service=Topic,name=PEStestTopic">
    <depends optional-attribute-name="DestinationManager">jboss.mq:service=DestinationManager</depends>
  </mbean>
</server>
```

```
</mbean>
```

```
</server>
```

This file must be placed in the directory `%JBOSS_HOME%/server/default/deploy`, and repository must be restarted. The topic will be bound to the jndi name `topic/PEStestTopic`.

JBoss comes with preconfigured queues and topics. These settings can be found in the file `jbossmq-destinations-service.xml`, which is located in the directory `%JBOSS_HOME%/server/default/deploy/jms`.

PASW Collaboration and Deployment Services Setup

To configure PASW Collaboration and Deployment Services to communicate with the JMS service:

1. Using Deployment Manager, set up a message domain referencing `PEStestTopic`.
2. Create a message-based schedule for an PASW Collaboration and Deployment Services job that uses the domain. For more information, see the Deployment Manager User's Guide.

Sample JMS Client Program

`SendMessageClientMain.java` is a sample program that publishes messages to the `topic/PEStestTopic` to subsequently trigger PASW Collaboration and Deployment Services processing.

```
import java.util.Properties;

import javax.jms.Connection;
import javax.jms.MessageProducer;
import javax.jms.Session;
import javax.jms.Topic;
import javax.jms.TopicConnectionFactory;
import javax.jms.TopicSession;
import javax.jms.TextMessage;
import javax.naming.InitialContext;
import javax.naming.Context;

public class SendMessageClientMain
{
    public static void main(String args[])
        throws Exception
    {
        System.out.println("SendMessageClient...");
        Connection connection = null;
        Session session = null;
        Topic topic = null;

        // Change these two lines to point to the correct topic and message to send
        String topicName = "topic/PEStestTopic";
```

```
String messageText = "testmsg";

Properties props = new Properties();
// Change these two lines to point the correct JMS server.
props.put(Context.INITIAL_CONTEXT_FACTORY, "org.jnp.interfaces.NamingContextFactory");
props.put(Context.PROVIDER_URL, "jnp://localhost:1099");

try {
    InitialContext iniCtx = new InitialContext(props);
    Object tmp = iniCtx.lookup("ConnectionFactory");
    TopicConnectionFactory tcf = (TopicConnectionFactory) tmp;

    connection = tcf.createConnection();
    topic = (Topic) iniCtx.lookup(topicName);
    session = connection.createSession(false,
        TopicSession.AUTO_ACKNOWLEDGE);
    connection.start();

    // Send a text msg
    MessageProducer send = session.createProducer(topic);
    TextMessage tm = session.createTextMessage(messageText);
    send.send(tm);
    System.out.println("sendMessage - sent text=" + tm.getText());
    send.close();
    System.out.println("End TopicSendClient");

    connection.stop();
    session.close();
    connection.close();
}
catch (Exception e)
{
    // Error handling left as an exercise for the student.
    e.printStackTrace();
}

System.exit(0);
}
```

The message topic, message text, and error handling routines can be modified as necessary. Instead of setting `Context.INITIAL_CONTEXT_FACTORY` and `Context.PROVIDER_URL` properties in the source code, the file *jndi.properties* defining how to look up the JMS server can be created as follows:

```
java.naming.provider.url=jnp://localhost:1099
java.naming.factory.initial=org.jnp.interfaces.NamingContextFactory
```

The file must be placed in the path of the compiled class of the sample program.

Running the Sample Program

The following is sample Windows batch file to publish a JMS message to JBoss Messaging service by running *SendMessageClientMain.class*. Note that the batch file sets environment variables to point to the appropriate JMS naming factory class prior to launching the sample program.

```
set JBOSS_HOME=C:\ssk\pgms\pes35\jboss\4.0.3SP1

set JBOSS_SERVER_LIB=%JBOSS_HOME%\server\default\lib
set JBOSS_ROOT_LIB=%JBOSS_HOME%\lib

set CLASSPATH=.\bin
set CLASSPATH=%CLASSPATH%;%JBOSS_SERVER_LIB%\jboss-j2ee.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_SERVER_LIB%\log4j.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_SERVER_LIB%\jnpserver.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_ROOT_LIB%\jboss-common.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_SERVER_LIB%\jbossmq.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_SERVER_LIB%\jboss.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_ROOT_LIB%\jboss-system.jar
set CLASSPATH=%CLASSPATH%;%JBOSS_ROOT_LIB%\concurrent.jar

java SendMessageClientMain
```

If all preceding steps have been completed correctly, running *SendMessageClientMain.class* will trigger the processing of the corresponding PASW Collaboration and Deployment Services message-based job schedule.

Message-Based Processing Example

Message-based scheduling functionality of PASW Collaboration and Deployment Services can be used to trigger processing by repository events as well as by third-party applications. For example, a job can be configured to be rerun when the PASW Modeler stream used in one of the job steps is updated. The procedure involves the following steps:

- ▶ Using Deployment Manager, create a JMS message domain.
- ▶ Set up a message-based schedule for the job using the message domain. Note that the JMS message selector must indicate the resource ID of the PASW Modeler stream as in the following example:

```
ResourceID=<resource ID>
```

The repository resource ID of the PASW Modeler stream can be found in the object properties.

- ▶ Set up a notification for the PASW Modeler stream based on the JMS subscriber you have defined.
- ▶ To test the message-based schedule, the stream must be opened in PASW Modeler, modified, and stored in repository. If everything has been set up correctly, the schedule will trigger the job.

Auditing the Repository

As the body of collected and created data objects grows, it is necessary to track the behavior of the data. Database auditing allows you to track the who, what, when, and how of data objects—who interacted with the data, what data objects were accessed, when the action took place, and how those objects were manipulated.

Depending on what level of detail is needed, repository provides a convenient mechanism for answering these questions, with the flexibility to gather as much or as little detail as required. Database reports and audits can be kept simple at first and become more complex as business needs change.

Note: On a day-to-day basis, changes to repository objects and processing results can be tracked through notifications and subscriptions. For more information, see the Deployment Manager documentation.

The practice of database auditing and reporting provides a way to:

- Monitor changes, such as the creation and removal of any data objects stored in the database.
- Record or log this database activity for future analysis and reference.
- Generate reports on database activity.

Being able to easily track these actions gives the user increased control over data and ensures compliance with the organization's rules for data security and change tracking.

Database Audit Facilities

repository provides several database tables for recording system events and changes to objects. When repository is installed in a supported relational database, the tables necessary for auditing and reporting are automatically created. The user is not required to populate any database objects manually

The easiest way to access auditing information is to run SQL queries in a supported database client application. For example, PASW BIRT Report Designer, included in the PASW Collaboration and Deployment Services installation, can be used to create auditing reports.

If certain kinds of auditing information must be retrieved on a regular basis, views can be set up. A database view is a read-only virtual or logical table composed of the result set of a query. Unlike ordinary tables in a relational database, a view is not part of the physical schema; it is a dynamic table computed or collated from data in the database. Changing the data in a table alters the data shown in the view.

repository is installed with several predefined views that can be used to retrieve a variety of auditing information about repository objects, including files, jobs, streams, etc. Custom views can also set up to meet more complex reporting requirements. When implementing custom views, refer to the database vendor's original documentation for variances in SQL syntax.

Note: Audit queries can be run against PASW Collaboration and Deployment Services event tables as well as the predefined views. However, because table structure may change in subsequent system releases, for compatibility considerations it is recommended to use views rather than tables when writing audit queries.

Audit Events

The following system events trigger entries into the database event tables:

repository Events

- Creating a file or folder
- Updating a file or folder
- Version
- Deleting a file or folder
- Modifying the permissions of a file or folder

Security Events

- Successful login
- Failed login
- Adding a user
- Deleting a user
- Changing a password
- Adding a group
- Adding a user to a group
- Deleting a group

Job Execution Events

- Submitting a job
- Starting a job
- Starting a job step
- Job successfully completes
- Job fails
- Job step successful
- Job step failure

Scoring Service Events

- Scoring request
- Scoring configuration change

Event Tables

repository event information is stored in audit event (SPSSAUDIT_EVENTS) and event parameter (SPSSAUDIT_PARAMETERS) tables. Every system event generates a row in the SPSSAUDIT_EVENTS table. An event can have associated parameter rows in the SPSSAUDIT_PARAMETERS table (one-to-many relationship).

Audit Events Table (SPSSAUDIT_EVENTS)

SERIAL. The unique identifier of the event row. The number can be used to determine the order in which the events were generated.

STAMP. The date and time when the event occurred.

COMPONENT. The system component originating the event. The following values may be returned for COMPONENT:

- repository/audit_component_name—Repository event
- security/componentAuthN—User authentication event
- security/componentLRU—User and group setup event
- prms/prms—Job scheduling event
- notification/notification—Notification or subscription event
- userpref/auditComponent—User preference change event
- scoring/scoring—Scoring service event

LOCUS. Defined by the owner component, assigns a more specific event type. The following values may be returned for LOCUS:

Repository Event Locus Codes

- repository/audit_access_object—File or folder accessed
- repository/audit_new_object—File or folder created
- repository/audit_update_object—File or folder updated (content or metadata)
- repository/audit_new_version—A version created
- repository/audit_delete_version—A version deleted
- repository/audit_delete_object—File or folder deleted
- repository/audit_move_object—File or folder moved
- repository/audit_modify_permissions—Permissions to a file or folder modified
- repository/audit_update_custom_property_value—Custom property value of a file or folder updated
- repository/audit_new_custom_property—New custom property created

- repository/audit_modify_custom_property—Existing custom property modified
- repository/audit_delete_custom_property—Existing custom property deleted
- repository/audit_reindex_repository_started—Repository re-index process started
- repository/audit_reindex_repository_ended—Repository re-index process ended

Security Event Locus Codes

- security/locAuthen—Successful login
- security/locNotAuthen—Failed login
- security/locLogout—Logout
- security/locLRUAdd—User added
- security/locLRUDelete—User deleted
- security/locLRUUpdate—Password change
- security/locLRUAdd—Group added
- security/locLRUUpdate—Group renamed
- security/locLRUUpdate—User added to/deleted from a group
- security/locLRUDelete—Group deleted

Job Execution Event Locus Codes

- prms/audit_job_submit—Job submitted
- prms/audit_job_start—Job started
- prms/audit_job_step_start—Starting a job step
- prms/audit_job_success—Job successfully completes
- prms/audit_job_failure—Job fails
- prms/audit_job_step_success—Job step successfully completes
- prms/audit_job_step_failure—Job step failure
- prms/audit_job_update—Job updated

Notification Event Locus Codes

- notification/audit_delivery—Notification message delivery event (delivered, not delivered, or partially delivered)
- notification/audit_subscription—Notifications or subscriptions settings change event (subscription created, updated, or deleted)

User Preference Event Locus Codes

- userpref/auditLSet—User preference value set
- userpref/auditLDelete—User preference value deleted

Scoring Service Event Locus Codes

- scoring/metric_update—Scoring service request or scoring configuration update

MIMETYPE. MIME type of the object associated with the event.

TITLE. Brief description of the event, generally shown in lists of events. For content repository events, this is the name of the file.

PRINCIPALID. The user that generated the event.

AUDIT_RESOURCE. If associated with content, this is the URI of the content repository object.

DETAILS. A string providing additional component-defined information about the event, such as the old label for label change, old metadata for metadata change, and the old name for name change.

SIGNATURE. Signature used to confirm the validity of data.

ADDRESS. The IP address of the client system associated with the event.

Audit Event Parameters Table (SPSSAUDIT_PARAMETERS)

SERIAL. The foreign key to the SPSSAUDIT_EVENTS table associating the parameter with the event.

NAME. A descriptive name of the parameter—for example, JobExecutionID, JobID, JobStepID, JobName, JobStepName, etc.

VALUE. The value of the named parameter.

Use database client application tools to obtain additional information about event table properties, such as column data types and nullability.

Audit Views

The following are audit views created in the database by default when repository is installed. Use database client application tools to obtain additional information about the properties of the views. Auditing database objects is performed by running SQL queries against the views. Note that the repository database also includes a number of other views that are used to support audit views. The support views are not intended for reporting.

Audit (SPSSPLAT_V_AUDIT)

The Audit view contains the auditing information from the File Version view. This view contains one row for every audit parameter for every audit event.

AUDITSERIALNUMBER. The unique identifier of the event. The number can be used to determine the order in which the events were generated.

AUDITTIMESTAMP. The timestamp of the audit (or the date an event was created) is set by the generating component.

AUDITCOMPONENT. The component or subsystem name that created the event and is under audit. The format is in the form `com.spss.<component>`.

AUDITCATEGORY. The category of events under audit.

MIMETYPE. The MIME type of the object under audit.

AUDITTITLE. The category or object name under audit.

AUDITPRINCIPAL. The principal user of object under audit.

AUDITRESOURCE. The contents host under audit, such as the content repository resource ID.

AUDITDETAILS. A string providing additional component-defined information about the event, such as the old label for label change, old metadata for metadata change, and the old name for name change.

ADDRESS. The IP address of the client system associated with the event.

AUDITPARAMETERNAME. An extended parameter of the audit event—for example, JobStepExecutionID, JobExecutionID, or JobID.

AUDITPARAMETERVALUE. An extended parameter value of the audit event—for example, the ID value.

AUDITRESOURCEID The repository ID of the resource associated with the event. Foreign key to the file or job ID in the File Version (SPSSPLAT_V_FILEVERSION) view.

AUDITMARKER Resource version associated with the event. Foreign key to the file or job version marker in the File Version (SPSSPLAT_V_FILEVERSION) view.

Custom Property (SPSSPLAT_V_CUSTOMPROPERTY)

The Custom Property view presents the file custom property information for the rows in the File Version view (one-to-many relationship).

PROPERTYNAME. The name of the custom property.

PROPERTYVALUE. The value of the custom property.

FILEID. Foreign key to the file or job in the File Version view to which this property applies.

File Version (SPSSPLAT_V_FILEVERSION)

The File Version view presents file and version information for repository objects such as PASW Modeler streams, PASW Statistics syntax files, SAS syntax files, etc. This view contains a row for every version of every file, folder, or job.

FILEID. The unique identifier of the file.

VERSION. The version of the file.

FILENAME. The name of the file.

VERSIONMARKER. The version marker for the file version.

VERSIONLABEL. The version label of the file version.

FILEPATH. The path to the file.

MIMETYPE. The mime type of the file.

AUTHOR. The author (user-specified) of the file.

DESCRIPTION. The description of the file.

FILECREATEDDATE. The date and time when the file was created.

FILECREATEDBY. The user who created the file.

FILELASTMODIFIEDDATE. The date and time when file was last modified.

FILELASTMODIFIEDBY. The user who last modified the file.

VERSIONCREATEDDATE. The date and time when the file version was created.

VERSIONCREATEDBY. The user who created the version of the file.

VERSIONLASTMODIFIEDDATE. The date and time when the file version was last modified.

VERSIONLASTMODIFIEDBY. The user who last modified the version.

Job History (SPSSPLAT_V_JOBHISTORY)

The Job History view presents job step execution information. This view contains a row for every execution for every job step in every job.

EXECUTIONID. The unique identifier of the execution.

JOBID. Foreign key to the job (FILEID) in the File Version view.

JOBVERSION. Foreign key to the job version in the File Version view.

JOBSTEPID. Foreign key to the job step in the Job Step view.

JOBSTEPEXECUTIONSTATUS. The success/failure status of the job step.

JOBSTEPEXECUTIONSTARTED. The start time of the job step.

JOBSTEPEXECUTIONENDED. The end time of the job step.

JOBSTEPEXECUTIONRUNTIME. The total run time of the job step.

JOBSTEPERRORLOG. The ID of the error log file for the job step.

JOBEXECUTIONSTATUS. The success/failure status of the job. The following values may be returned for **JOBEXECUTIONSTATUS**:

- Null—Unknown
- 0—Failure
- 1—Success
- 2—Queued
- 3—Running
- 4—Ended
- 5—Cascading
- 6—Error
- 7—Cascade error
- 8—Canceling
- 9—Canceled
- 10—Cancel pending
- 11—Cascade canceled
- 12—Joining

JOBEXECUTIONSTARTED. The start time of the job.

JOBEXECUTIONENDED. The end time of the job.

JOBEXECUTIONRUNTIME. The total run time of the job.

JOBCLUSTERQUEUEDDATETIME. The time the job was placed in the queue. The job queued time is slightly later than the submitted time.

JOBCLUSTERCOMPLETIONCODE. Depending on job type, this is an integer value that corresponds to the job status. Zero (0) indicates success for all types of jobs.

JOBCLUSTERAPPLICATIONSTATUS. Depending on job type, this is a string value that corresponds to the job status.

JOBPROCESSID. Depending on the type of job, this is the ID of the corresponding system process—for example, the operating system process ID for a running executable file.

JOBEXECUTEDPARAMETERS. This field currently is not being used.

JOBNOTIFICATIONENABLED. Indicates whether notification is enabled for the job.

Job Step (SPSSPLAT_V_JOBSTEP)

The Job Steps view contains the information about job steps in jobs. This view contains a row for every job step for every version of every job.

JOBSTEPID. The unique identifier of the job step.

JOBSTEPNAME. The name of the job step.

JOBID. Foreign key to the job (FILEID) in the File Version view containing this job step.

JOBVERSION. Foreign key to the job version in the File Version view containing this job step.

JOBSTEPTYPE. The type of the job step. Currently, the types include ClementineStreamWork, SPSSSyntaxWork, SASSyntaxWork, ExecutableContentWork (General Work), and WindowsCommandWork. Related DOS commands can be either of WindowsCommandWork or ExecutableContentWork type.

REFERENCEDFILEID. The ID of the file referenced by this job step, if applicable—for example, a PASW Modeler stream, an PASW Statistics or SAS syntax file, etc.

REFERENCEDFILELABEL. The label of the file referenced by this job step, if applicable.

Schedule (SPSSPLAT_V_SCHEDULE)

The Schedule view presents the schedule information that is associated with a job in the File Version view. This view contains a row for every schedule.

JOBID. Foreign key to the job (FILEID) in the File Version view.

JOBVERSION. Foreign key to the job version in the File Version view. This is the version of the job to execute at this time. If the job label is moved (or if a new job version is saved and the schedule is set to execute the latest job), the job version will change.

SCHEDULEDFREQUENCY. The schedule recurrence relates to the scheduled interval and time units. For example, if frequency is daily and interval is 1, then scheduled day of week can be any day from Sunday to Saturday, while scheduled day of month will be 0.

SCHEDULEDINTERVAL. This is the number of intervals to skip between schedules. The meaning changes based on the value of SCHEDULEDFREQUENCY—for example, a frequency of weekly and an interval of 4 means run every fourth week.

SCHEDULEDDAYOFMONTH. The day of the month for monthly schedules.

SCHEDULEDDAYOFWEEK. The day of the week for weekly schedules.

SCHEDULEDTIME. The scheduled time that the job will start.

SCHEDULESTARTDATE. The start date for recurring schedules (daily, weekly, monthly), or the date to execute for other schedules.

SCHEDULEENDDATE. The end of recurrence date for the recurring schedules of type daily, weekly, monthly. This column will be null for the other schedule types, and may be null for the listed schedule types if the schedule is to stop triggering at the listed date.

NEXTSCHEDULED TIME. The next start date of the schedule. It will be null if the schedule is past its end date or is a one-time schedule.

SCHEDULEENABLED. Schedule enabled.

SCHEDULELABEL. The label of the job to execute when the schedule triggers.

SCHEDULELASTUPDATE. The date timestamp that this schedule was last modified.

SCHEDULECREATOR. The user ID of the person who created the schedule.

Stream Attribute Value (SPSSPLAT_V_STREAMATTRVALUE)

The Stream Attribute Value view presents the attribute information about the nodes in a PASW Modeler stream. This view contains a row for every allowable value of every attribute in every stream.

ATTRIBUTEID. The unique identifier of the attribute.

ATTRIBUTENAME. The name of the attribute.

NODEID. Foreign key to the node in the Stream Node view.

ATTRIBUTETYPE. The attribute type.

ATTRIBUTE CATEGORICAL VALUE. An allowable value for the attribute for multivalued attributes.

NUMERICAL UPPER BOUND. The upper bounds value allowable for numerical attributes.

NUMERICAL LOWER BOUND. The lower bounds value allowable for numerical attributes.

Stream Node (SPSSPLAT_V_STREAMNODE)

The Stream Node view presents the information for the nodes in PASW Modeler streams. This view contains a row for every node in every version of every stream.

NODEID. The unique identifier of the node in the stream.

STREAMID. Foreign key to the stream (FILEID) in the File Version view containing this node.

STREAMVERSION. Foreign key to the stream version in the File Version view containing this node.

NODENAME. The name of the node in the stream.

NODETYPE. The type of the node in the stream.

NODELABEL. The label of the node in the stream.

ALGORITHMNAME. The algorithm of the node for modeling nodes.

MININGFUNCTION. The data mining function of the node for modeling nodes.

IOFILENAME. The input or output file of the node, for FileInput or FileOutput nodes.

IODATABASETABLE. The name of the database table name for DatabaseInput or DatabaseOutput nodes.

IODSN. The data source name of the node for DatabaseInput or DatabaseOutput nodes.

Note: For this release, the ioDSN column in the SPSSPLAT_V_STREAMNODE view is not used. This column will contain NULL for each record.

Scoring Service Logging

PASW Collaboration and Deployment Services also provides database facilities for logging the operations of the scoring services. The following database objects are used to store the scoring service information:

- Request log table
- Database views
- XML schema

Scoring service logging is supported on all database management systems that can be used for the repository:

- DB2
- MS SQL Server
- Oracle

Note: DB2 on IBM i cannot be used for scoring service logging.

Request Log Table

By default, the Scoring Service request information is stored in SPSSSCORE_LOG table.

Scoring Log Table (SPSSSCORE_LOG)

SERIAL. The unique identifier of the Scoring Service request.

STAMP. The date and time of the Scoring Service request.

INFO. Additional information about the scoring request in XML format. The information is generated according to the XML schema registered with the database. The same information is available in relational format from the scoring log view.

Cleanup and Maintenance

Over time, as Scoring Service requests are logged, the SPSSSCORE_LOG can become quite large and it may be necessary to delete records from this table. For example, the administrator may to purge old records prior to January 1, 2009 by running the following SQL statement:

```
DELETE FROM spssscore_log WHERE STAMP < '2009-01-01'
```

Database Views

The following scoring views are created in the database by default when the repository is installed. They present the information stored as XML in the INFO column of SPSSSCORE_LOG table in relational format. Use database client application tools to obtain additional information about the properties of the views or run SQL queries.

Scoring Request (SPSSSCORE_V_LOG_HEADER)

This view contains a row for every scoring request row in the SPSSSCORE_LOG table.

SERIAL. The unique identifier of the scoring request.

ADDRESS. The IP address for the machine initiating the scoring request. Note that in certain cases it may be the address of the server rather than the client, for example, the address of the cluster load balancer or proxy server.

HOSTNAME. The name of the machine initiating the scoring request. If the servlet container running the Scoring Service on this machine does not allow Domain Name System reverse lookups, the value corresponds to the IP address of the machine. If no host name can be determined, a null value is used. In cases when hostname lookup takes too long, it may be possible to improve Scoring Service performance by configuring the system not to look up the hostname using the corresponding configuration option in browser-based Deployment Manager.

PRINCIPAL. The user name associated with the scoring request. If this value is not included in the request, no information is logged.

STAMP. This column contains the timestamp of when the Scoring Service logged the request.

MODEL_OBJECT_ID. The repository identifier of the object that was configured with the Scoring Service. For example, if a PASW Modeler stream was configured for scoring, this is the repository identifier of the stream.

MODEL_VERSION_MARKER. The identifier of the specific version of the repository object that was configured for scoring.

CONFIGURATION_NAME The name of the Scoring Service configuration entry. The name is assigned when a model is configured for scoring.

Scoring Request Input (SPSSSCORE_V_LOG_INPUT)

The view contains the information about the model inputs that were used to produce the score. There may be multiple rows in SPSSSCORE_V_LOG_INPUT for each row in SPSSSCORE_LOG table and SPSSSCORE_V_LOG_HEADER view. Each row in the SPSSSCORE_V_LOG_HEADER represents a single input value.

SERIAL. The unique identifier of the scoring request row.

INPUT_TABLE. If the input source is the Enterprise View, this is the Enterprise View table name.

INPUT_NAME. The name of an input field. If the input source is the Enterprise View, this is the Enterprise View column name.

INPUT_VALUE. Input value.

INPUT_TYPE. Input data type. The following data types are allowed:

- date
- daytime
- decimal
- double
- float
- integer
- long
- string
- timestamp

Scoring Request Context Data (SPSSSCORE_V_LOG_CONTEXT_INPUT)

This view contains the information about the data that was passed to the Scoring Service and used as a Context Data Source for the Enterprise View Real-Time DPD. There may be multiple rows in SPSSSCORE_V_LOG_CONTEXT_INPUT view for each row in SPSSSCORE_V_LOG_HEADER view.

SERIAL. The unique identifier of the scoring request row.

CONTEXT_TABLE. The name of the table used in the Context Data Source.

CONTEXT_ROW. The row number of the context data row starting at 1.

CONTEXT_NAME. The name of an input field corresponding to the name of the column in the Context Data Source.

CONTEXT_VALUE. Input value.

Scoring Request Output (SPSSSCORE_V_LOG_OUTPUT)

The SPSSSCORE_V_LOG_OUTPUT view is used to log the outputs of the Scoring Service. There may be multiple rows in SPSSSCORE_V_LOG_OUTPUT view for each row in SPSSSCORE_V_LOG_HEADER view. The Scoring Service has the ability to provide multiple outputs. Each output can consist of multiple values. For example, the Scoring Service may provide two recommendations (two outputs). Each of these recommendation will be assigned a unique row number starting at 1. For each recommendation, there may be multiple output values.

SERIAL. The unique identifier of the scoring request row.

OUTPUT_ROW. The row number of context data row starting at 1.

OUTPUT_NAME. The output field name (attribute name) corresponding to the name of the column in the Context Data Source.

OUTPUT_VALUE. Output value.

Scoring Request Metrics (SPSSSCORE_V_LOG_METRIC)

The SPSSSCORE_V_LOG_METRIC view is used to log the output metrics of the Scoring Service, for example, the time to process the scoring request. There may be multiple rows in SPSSSCORE_V_LOG_METRIC view for each row in SPSSSCORE_V_LOG_HEADER view.

SERIAL. The unique identifier of the scoring request row.

METRIC_NAME. The name of an metric field.

METRIC_VALUE. Metric value.

Scoring Request Properties (SPSSSCORE_V_LOG_PROPERTY)

The SPSSSCORE_V_LOG_PROPERTY view is used to log the properties used in processing the request. There may be multiple rows SPSSSCORE_V_LOG_PROPERTY view for each row in SPSSSCORE_V_LOG_HEADER view. The properties that can be logged depend on the selected score provider.

SERIAL. The unique identifier of the scoring request row.

METRIC_NAME. The name of a property.

OUTPUT_VALUE. Property value.

Audit Query Examples

The following are examples of SQL queries against audit views. Note that certain SQL functions are specific to Microsoft SQLServer and may be invalid on other database platforms.

Successful login attempts for user 'jsmith'

```
select AUDITTIMESTAMP as "Login date",  
ADDRESS as "Machine address"  
from SPSSPLAT_V_AUDIT  
where AUDITCOMPONENT = 'security/componentAuthN'  
and AUDITCATEGORY = 'security/locAuthen'  
and AUDITTITLE = 'jsmith'  
order by 1 desc
```

Unsuccessful login attempts for all users

```
select AUDITTITLE as "Username",  
AUDITTIMESTAMP as "Login date",  
ADDRESS as "Machine address"  
from  
SPSSPLAT_V_AUDIT  
where AUDITCOMPONENT = 'security/componentAuthN'  
and AUDITCATEGORY = 'security/locNotAuthen'  
order by 1 asc, 2 desc
```

Number of successful login attempts for each user over the last month

```
select AUDITTITLE as "Username",  
COUNT(*) as "Successful logins"  
from  
SPSSPLAT_V_AUDIT  
where AUDITCOMPONENT = 'security/componentAuthN'  
and AUDITCATEGORY = 'security/locAuthen'  
and AUDITTIMESTAMP >= DATEADD(month, -1, GETDATE())  
group by AUDITTITLE  
order by 2 desc
```

All repository resources that have custom property 'Region'

```
select V1.FILEPATH + V1.FILENAME as "Resource", V2.PROPERTYNAME + ' = ' + V2.PROPERTYVALUE as "Property/Value"  
from SPSSPLAT_V_FILEINFO V1,  
SPSSPLAT_V_CUSTOMPROPERTY V2  
where V1.FILEID = V2.FILEID  
and V2.PROPERTYNAME = 'Region'
```

All repository resources that have custom property value 'Asia-Pacific'

```
select V1.FILEPATH + V1.FILENAME as "Resource", V2.PROPERTYNAME + ' = ' + V2.PROPERTYVALUE as "Property/Value"  
from SPSSPLAT_V_FILEINFO V1,  
SPSSPLAT_V_CUSTOMPROPERTY V2  
where V1.FILEID = V2.FILEID  
and V2.PROPERTYVALUE = 'Asia-Pacific'
```

All repository resources modified (new versions created) by user 'jsmith'

```
select FILEPATH + '/' + FILENAME as "Resource",  
VERSION as "Version",  
VERSIONCREATEDDATE as "Modified date"  
from SPSSPLAT_V_FILEVERSION  
where VERSIONCREATEDBY = 'jsmith'
```

All users who modified file /Modeler/Base_Module/drugplot.str

```
select VERSION as "Version",  
VERSIONCREATEDBY as "Username",  
VERSIONCREATEDDATE as "Created date"  
from SPSSPLAT_V_FILEVERSION  
where FILEPATH + FILENAME = '/Modeler/Base_Module/drugplot'
```


Troubleshooting

Certain error messages and symptoms are common when installing and working with PASW Collaboration and Deployment Services. Methods for clearing these errors and establishing a functional system exist for:

- **PASW Collaboration and Deployment Services.** Common problems when installing and starting the application on supported server platforms.
- **Solaris 9.** Known issues related to PASW Collaboration and Deployment Services on Sun's UNIX operating system.
- **HP-UX.** Known issues related to PASW Collaboration and Deployment Services on HP UNIX operating system.
- **DB2 for IBM i.** Symptoms and error messages that surface while transacting with a DB2 database running on IBM i.
- **Oracle 10g and 11g.** Symptoms and error messages that surface while transacting with an Oracle 10g and 11g databases.
- **JBoss.** JBoss application server running PASW Collaboration and Deployment Services.
- **Oracle WebLogic.** WebLogic application server running PASW Collaboration and Deployment Services.
- **WebSphere.** WebSphere application server running PASW Collaboration and Deployment Services.

It is always a good practice to refer to PASW Collaboration and Deployment Services log files to establish the cause of the problem.

PASW Collaboration and Deployment Services

How do I prevent performance bottlenecks and CPU usage issues when starting and deploying PASW Collaboration and Deployment Services?

Depending on the specific system configuration, previously installed antivirus or spyware software may be configured for “deep scanning” of application components. These third party applications can be reconfigured to scan during certain times, or they can be turned off during installation and manually restarted.

Additionally, some of the more strict server-side firewall settings may negatively impact startup performance and not allow access.

If you are experiencing significant system degradation when starting the service, disable any nonessential processes and restart the PASW Collaboration and Deployment Services.

Once I log in to the administrative interface, how do I determine which database I am accessing?

Database connection information can be downloaded and accessed from the Web interface.

1. After authenticating, click About from the navigation list options. The About page appears.
2. Click the Download version and system details link at the bottom of the page. When prompted, save the file to disk.
3. Open the file in a text editor and search for *Database Details*. This section contains detailed information on the database being used, including name, version, and a table listing.

The application throws java.lang.OutOfMemoryError: PermGen space exception.

This error occurs when the JVM runs out of space in the permanent generation heap due to a large number of used classes. The solution is to increase the value specified with PermSize JVM parameter. For example, for JBoss installations, the size of permanent generation heap available to the wrapper service can be increased by modifying the following line in *<JBoss Installation Directory>/wrapper/conf/wrapper.conf*:

```
wrapper.java.additional.1=-Dprogram.name=run.bat -XX:PermSize=128m.
```

For information about increasing the permanent generation heap size for other application servers, see the application server vendor documentation.

Out of memory errors can also be prevented by adding JVM parameters to tune memory allocation and garbage collection, for example:

```
-XX:+CMSPermGenSweepingEnabled -XX:+CMSClassUnloadingEnabled
```

When a BIRT report is run in Deployment Portal, the application is not able to authenticate my credential for accessing the data source of the report and is repeatedly displaying the login screen.

- Verify that the data source for the report and the credentials are defined correctly. For more information, see the corresponding section of the *Deployment Manager User's Guide*.
- If the data source for the report is JDBC-based, verify that the proper driver is installed with repository. For driver path information specific to the operating platform, see the installation instructions.

SAS syntax job processed in PASW Collaboration and Deployment Services running on a UNIX system fails with to a database connection error due to invalid library name ("ERROR: Error in the LIBNAME statement").

- Verify that the shared libraries path environment variable (LD_LIBRARY_PATH on Solaris, SHLIB_PATH on HP-UX, or LIBPATH on AIX) is set to an appropriate value.

How do I restore PASW Collaboration and Deployment Services if my keystore file has been lost?

The keystore file contains the keys used to encrypt passwords used by PASW Collaboration and Deployment Services, such as the master password for database access. If the keystore file is lost, the system becomes unusable. If backup of the keystore is available, it can be restored to the

original location. If you are unsure what the original path of the keystore was, you can look up the *keystorePath* property of *keystoreSecurity* element in *<PASW Collaboration and Deployment Services Installation Directory>/platform/setupinfo.xml*.

If the keystore file is lost and backup is not available, the system must be reinstalled by re-running the setup utility in *<PASW Collaboration and Deployment Services Installation Directory>/setup* and pointing it to the existing repository database. All passwords that existed in the system, such as the passwords for external directory services, defined credentials, etc. must be manually reentered.

A BIRT report against DB2 IBM i V6R1 database using prompted credentials fails when run in PASW Collaboration and Deployment Services .

Add `prompt=true` parameter to the JDBC connection URL.

```
Driver Name: com.ibm.as400.access.AS400JDBCdriver
Driver URL: jdbc:as400://myServer/B101E31E;prompt=false
```

“Build New Scoring Configuration Details Failed” error when configuring scoring on non-Windows PASW Collaboration and Deployment Services installations

“Build New Scoring Configuration Details Failed” error message is displayed when scoring configuration dialogue is opened in Deployment Manager. The problem is corrected by changing the permissions on *<PASW Collaboration and Deployment Services installation directory>/components/modeler/modelerserver* file to `execute`, for example:

```
cd /usr/PASWCDS4/components/modeler/modelerserver
sudo chmod +x modelerserver
```

Reporting output generated as a PDF file does not display national character sets correctly

On certain UNIX systems, the default JVM font configuration may not be suitable for all national character sets, such as Asian language characters. In these cases, it may be necessary to specify the default JVM font using a font configuration file. For information about Java font configuration files, see Sun documentation.

Solaris

How do I avoid getting an access error message when trying to run the installation script?

PASW Collaboration and Deployment Services must be installed by a user with adequate privileges. Change the active user to *root* (or to another user with adequate access rights) and run the installation script.

To which other directories does the installing user need access?

The user running the installation must also have write access to */etc/.java* for the system to function properly.

If the installation is executed by a user without write access to `/etc/.java`, switch to a user with write access and run the setup shell script again. Once the installation is complete, verify that the following file exists:

```
/etc/.java/.systemPrefs/com/spss/setup/component/services/prefs.xml
```

Unable to start PASW Collaboration and Deployment Services on JBoss and Solaris 9.

When attempting to start PASW Collaboration and Deployment Services on JBoss and Solaris 9, “`ld.so.1: wrapper: fatal: libm.so.2: open failed: No such file..`” error occurs.

To resolve the problem, create symbolic link `/usr/lib/64/libm.so.2` to `/usr/lib/64/libm.so.1`:

```
ln -s /usr/lib/64/libm.so.1 /usr/lib/64/libm.so.2
```

HP-UX

Import failure when running PASW Collaboration and Deployment Services on HP-UX with NFS.

When importing resources into the PASW Collaboration and Deployment Services repository running on HP-UX with NFS, the following exception may occur:

```
java.lang.RuntimeException: The database is already in use by another process: org.hsqldb.persist.NIOLockFile@3ffdc36b[file
=/qa/projects/pes/HPUX/appserv/bea11g/user_projects/domains/PASWDomain41B179a/pasw_transfer_root/
0a0b0ad397fef2c500000126b4ca991881ab/0a0b0ad397fef2c500000126b4ca991881ad_transfer_database.lck,
exists=true, locked=false, valid=false, fl=null]:
```

To resolve the problem, use browser-based Deployment Manager to set the value of *Repository* -> *Resource Transfer Lookup Table* configuration option to **MEMORY**. For more information, see PASW Collaboration and Deployment Services administrator’s documentation.

Oracle Database

How do I create a user and tablespace?

To clear and reestablish the *spssplat* user and tablespace from an Oracle database, issue the following set of commands:

```
drop user spssplat cascade; CREATE USER spssplat IDENTIFIED BY spssplat
DEFAULT TABLESPACE SPSSPLAT TEMPORARY TABLESPACE TEMP
QUOTA UNLIMITED ON SPSSPLAT;
@$ORACLE_HOME/sqlplus/admin/pupbld;
GRANT CONNECT, RESOURCE, UNLIMITED TABLESPACE TO spssplat;
```

JBoss

How is the session timeout value configured to adjust the amount of time a user can remain idle?

Once a user is logged in to PASW Collaboration and Deployment Services, a period of inactivity is allowed before the session is terminated and the user must reauthenticate. To increase or decrease this value:

1. From the installation directory, navigate to `\JBoss\server\default\deploy\jbossweb-tomcat50.sar\`.
2. Open `web.xml` in a text editor.
3. Locate the section for *Default Session Configuration*, and edit the value for `<session-timeout>`.
4. Stop and restart the application.

Note: This file is processed when the application is deployed; configuration changes do not take effect until the server is restarted.

How do I determine the port on which my version of JBoss is running?

The JBoss application server's HTTP port is defined in the file:

```
jboss-3.2.7\server\default\deploy\jbossweb-tomcat50.sar\server.xml
```

with the attribute:

```
/Server/Service/Connector@port
```

Note: Depending on the release of JBoss, the version numbers in the path may vary.

What additional settings are required for PASW Collaboration and Deployment Services FIPS 140-2 compliance on JBoss?

For PASW Collaboration and Deployment Services to function properly when running on JBoss in FIPS 140-2-compliant mode, `{URIEncoding="UTF-8"}` attribute must be specified for the HTTPS connector.

Alternatively, from the command line, the `netstat` command can be used to view applications and the ports that are in use.

WebLogic

"IOException: Resource has been deleted" is thrown in Deployment Portal when trying to access file attachments that contain reporting output.

The exception can occur if the PASW Collaboration and Deployment Services installation is running on WebLogic application server using JRockit rather than Sun JRE. If the exception occurs, reconfigure WebLogic to use Sun JRE. For more information, see WebLogic documentation.

Cascading parameters are not displayed correctly in reports when PASW Collaboration and Deployment Services is run with WebLogic 9.2 and 10 on Solaris 10.

-Djava.awt.headless=true startup argument must be added to the application server Java environment.

PASW Collaboration and Deployment Services setup on Red Hat v5.4 fails with “Too many files open” message.

This error is generated when the open file limit for a user exceeds the default setting. You can check the user’s open file limit with the following command:

```
ulimit -n
```

The user’s open file limit can be increased by editing */etc/security/limits.conf*, for example, appending the following line:

```
@username - nofile 2048
```

The system must be restarted for the new limit to take effect.

WebSphere

Miscellaneous errors occur during package installation (with Package Manager) into the repository using a WebSphere application server.

Make sure the latest vendor patches have been applied to the application server.

Server log is reporting encryption errors, such as exception com.ibm.crypto.provider.AESCipher.engineGetKeySize(Unknown Source)

The error occurs with WebSphere 6.1 Service Pack 19 and is caused by the incorrect password value. To correct the error, copy the value of platform.keystore.password from

```
<PASW Collaboration and Deployment Services installation directory>/platform/setupinfo.xml
```

to

```
<WEBSHERE_HOME>/profiles/AppSrv01/config/cells/xi-wyueNode01Cell/nodes/xi-wyueNode01/servers/
<server name>/server.xml
```

Upgrading to WebSphere 6.1 Service Pack 23 may also resolve encryption problems.

“CWSIS1535E: The messaging engine’s unique id does not match that found in the data store” error

The error can be corrected by stopping PASW Collaboration and Deployment Services and deleting the repository database tables with names beginning with the *SIB* prefix. The tables will be recreated when PASW Collaboration and Deployment Services is restarted. Note that this solution applies only if you do not need to keep any of the currently stored

persistent messages. For more information about WebSphere JMS troubleshooting, see <http://www.redbooks.ibm.com/redpapers/pdfs/redp4076.pdf>.

Nativestore Schema Reference

The *nativestore.xsd* schema defines the structure of an XML file containing users and groups to be imported into the Deployment Manager. In addition, the file can specify obsolete users and groups that should be deleted.

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore>
  <user userID="sbennett" password="sb9482" encrypted="false">
    <group>sales</group>
  </user>
  <user userID="lsanborn" password="ls7725" encrypted="false">
    <group>sales</group>
  </user>
  <user userID="lalger" password="la4011" encrypted="false">
    <group>analyst</group>
  </user>
  <user userID="cjones" password="cj2683" encrypted="false">
    <group>analyst</group>
  </user>
  <obsolete>
    <user>mmonroe</user>
    <user>bgmurphy</user>
    <group>jones project</group>
  </obsolete>
</nativestore>
```

nativestore Element

Child elements: [user](#), [obsolete](#)

Root element for importing local users and their groups into the Deployment Manager.

user Element

Parent element: [nativestore](#)

Child elements: [group](#), [role](#)

User to be added or updated.

Table B-1
Attributes for the user element

Name	Type	Use	Default	Description
userID	string	required	<i>no default value</i>	User ID that will be used to log in to the system.
password	string	optional	<i>no default value</i>	Usually a plain-text password. If the <code>encrypted</code> attribute is true, then this password is encrypted. It is generally not practical to use an encrypted password when importing. Passwords are encrypted when exporting from the server, but this is <i>not</i> exposed in the Deployment Manager user interface.
encrypted	boolean	optional	false	Indicates if the password is plain-text or encrypted. Encrypted passwords are exported from the native store (encryption is one-way, making it impossible to re-create a user's password). When importing from another system, passwords must be plain-text; the <code>encrypted</code> attribute is usually omitted.

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore>
  <user userID="sbennett" password="sb9482" encrypted="false">
    <group>sales</group>
  </user>
</nativestore>
```

group Element

Type: string

Parent element: [user](#)

Groups associated with the user. If a group does not exist, it will be created automatically.

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore>
  <user userID="sbennett" password="sb9482" encrypted="false">
    <group>sales</group>
  </user>
</nativestore>
```

role Element

Type: string

Parent element: [user](#)

Role associated with the user. If a role does not exist, it will *not* be added automatically.

obsolete Element

Parent element: [nativestore](#)

Child elements: [user](#), [group](#)

Groups or users to be removed. Note that they may be loaded in “replace mode,” which will automatically remove all groups and non-administrative users. In that mode, this element has no effect.

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore>
  <obsolete>
    <user>mmonroe</user>
    <user>bgmurphy</user>
    <group>jones project</group>
  </obsolete>
</nativestore>
```

user Element

Type: string

Parent element: [obsolete](#)

The user ID to be removed. A user with administrative privileges cannot be removed.

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore>
  <obsolete>
    <user>mmonroe</user>
  </obsolete>
</nativestore>
```

group Element

Type: string

Parent element: [obsolete](#)

Group name to be removed.

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore>
  <obsolete>
    <group>jones project</group>
  </obsolete>
</nativestore>
```

Index

- actions, 22
 - adding to roles, 39
 - removing from roles, 39
 - roles, 36
- Active Directory, 22, 42, 49
 - disabling, 43
 - enabling, 43
 - with local override, 42, 44, 50
- Active Directory with Local Override, 22–23
- adding
 - administered servers, 16
 - groups, 29
 - MIME types, 73
 - users, 26
- administered servers
 - adding, 16
 - deleting, 20
 - logging in, 19
 - logging out, 20
 - properties, 19
 - server information, 17
 - types, 16
- administrative privileges, 55–56, 59–60
- administrators, 38
- allowed users, 22, 34
 - for Active Directory, 45
- Apache ActiveMQ, 93
- Asian languages, 115
- audit queries, 110
- audit reports, 97
- audit tables, 97
- audit views, 97
- auditing, 91, 97
 - database schema, 99
 - events, 98

- BEA WebLogic, 93
- BIRT, 64
- BIRT report processing, 114

- cache, 56
- caching
 - logins, 69
- capturing audit events, 98
- cascading stylesheets, 64
- changing
 - passwords, 13
- Coherence, 56
- collaboration, 1
- components, 14

- configuration , 55–60, 64, 66, 68–71
 - Deployment Portal scoring, 58
 - options, 87
 - scoring, 58
- configuring
 - cache, 56
 - custom dialogs, 57
 - Data Service, 57
 - Deployment Manager, 57
 - Deployment Portal, 58
 - Enterprise View, 59
 - Help, 55, 60
 - notification, 60
 - pager, 64
 - PASW Statistics, 57
 - process management, 64
 - repository, 66
 - security, 55, 69
 - setup, 70
 - ShowCase Warehouse Builder, 71
 - system, 55–60, 64, 68–71
 - templates, 55–56
 - URL prefix, 70
- console
 - UNIX, 8, 11
 - Windows, 8, 11
- conventions
 - naming, 21
- creating
 - allowed users, 34
 - extended groups, 32
 - groups, 29
 - roles, 38
 - users, 26
- credentials, 57
- custom dialogs, 57
- customizing
 - message templates, 77, 82
 - notification messages, 77, 82
 - notifications, 78, 80
- CWSIS1535E error, 118

- Data Service
 - configuration, 57
- database auditing, 97
- database lock exception, 116
- database schema
 - auditing, 99
- databases
 - troubleshooting, 116
- DB2, 115

- debugging information, 66
- debugging the notification service, 91
- dedicated SMTP server, 88
- deleting
 - administered servers, 20
 - groups, 32
 - MIME types, 74
 - users, 29
- delivery failure, 91
- delivery status notifications, 91
- deployment, 2
- Deployment Manager, 3–4
 - configuration, 57
- Deployment Portal, 3–4
 - configuration, 58
- Deployment Portal scoring configuration, 58
- diagnosing errors, 113, 115–116
- directory path, 66
- disabling binary content , 88
- domain, 52
- driver URL, 115
- DSN, 91

- e-mail notifications, 77
 - HTML, 82
 - text, 82
- editing
 - groups, 31
 - MIME types, 74
 - roles, 39
 - users, 28
- EIM, 42, 50
- eim.jar, 42
- encrypted attribute
 - for user, 121
- encryption, 114
- Enterprise Identity Management, 42, 50
- Enterprise View, 3, 5, 59
- environment variables, 114
- error messages, 113, 115–116
- errors, 113, 115–116
 - access, 115
 - diagnosing, 113, 115–116
 - generation heap size, 113
 - installation, 113
 - java.lang.OutOfMemoryError: PermGen space, 113
 - memory errors, 113
 - resolving, 113, 115–116
 - wrapper service, 113
- event collector pool, 88
- event noise filtering, 88
- events
 - auditing, 98
 - job execution, 98
 - repository, 98
 - security, 98
- execution servers, 5
 - remote process, 6
 - SAS, 6
- exporting, 38
- extended groups, 22, 32
 - for Active Directory, 45
- external security provider, 22
 - Active Directory, 22
 - Active Directory with Local Override, 22
 - OpenLDAP, 22

- file permissions, 115
- files
 - associating with images, 74
 - naming, 21
- FIPS 140-2
 - JBoss configuration, 117
- folders
 - naming, 21
- fonts, 115

- garbage collection, 114
- generation heap size, 113
- group element
 - in obsolete, 122
 - in user, 120–121
- groups
 - adding, 23, 29
 - creating, 23, 29
 - deleting, 32
 - editing, 23, 31
 - extended, 22–23, 32
 - importing, 32
 - local, 23
 - managing in Deployment Manager, 22
 - modifying, 23, 31
- guidelines
 - naming, 21

- heap size, 114
- Help, 55, 60
- HP-UX, 116

- IBM i, 115
- IBM i user repository, 42
- images
 - associating with files, 74
- import failure, 116
- importing, 38
- importing users and groups, 32
- indexing
 - authority to perform, 75
 - configuration option to force, 75
 - on repository upgrade, 75

- installation errors, 113
- installed packages, 14

- Java Messaging Service, 93–94, 96
- java.lang.OutOfMemoryError: PermGen space, 113
- jBoss, 91
- JBoss Messaging, 93
- JBoss naming factory, 93
- JDBC, 115
- JDBC drivers, 114
- JMS, 93–94, 96
- JMS bus, 118
- JMS message domain, 94, 96
- JMS queue, 93
- JMS topic configuration, 93
- JMS topics, 93
- JNDI, 93–94
- jndi.properties, 94
- job execution events, 98
- job output
 - viewing, 87
- job status, 84
- job step history, 85
- JobStatusURL property
 - in notification templates, 84

- Kerberos
 - domain, 52
 - JAAS, 52
 - Key Distribution Center, 52
 - key table file, 52
 - realm, 52
 - Service Ticket, 52
- keystore file, 114
- keystore file backup, 114

- LD_LIBRARY_PATH, 114
- LDAP, 42
- LIBPATH, 114
- license, 14
- local groups
 - for Active Directory, 50
- local override
 - for Active Directory, 42
- local principal filter
 - for Active Directory, 50
- local security provider, 22
- local user repository, 42
- login, 8
- login page, 12–13
- logins
 - caching, 69
- logout, 8
- logs, 14

- memory allocation, 114
- memory errors, 113–114
- message-based processing example, 96
- message-based scheduling, 93–94, 96
- messageContent element
 - contentType attribute, 82
 - in notification templates, 77, 80, 82
- messageProperty element
 - in notification templates, 77–78
- messageSubject element
 - in notification templates, 77, 80
- MIME, 72
- MIME types, 72, 82
 - adding, 73
 - deleting, 74
 - editing, 74
- mimeMessage element
 - in notification templates, 77
- missing JDBC drivers, 114
- modifying
 - groups, 31
 - users, 28

- naming conventions, 21
- national character sets, 115
- native provider, 42–43, 46, 48–50
- nativestore element, 120
- nativestore schema, 120
- navigation, 8, 13
- NFS, 116
- notification
 - configuration, 60
- notification configuration options, 87
- notification delivery failure, 91
- notification performance recommendations, 87
 - number of custom templates, 89
 - number of recipients, 89
 - number of subscriptions, 89
 - subscriptions management, 89
- notifications, 77
 - content, 77
 - customizing, 78, 80, 82
 - formatting, 82
 - HTML, 82
 - subject header, 77
 - templates, 77, 83
 - text, 82
 - Velocity, 77
- nslookup, 91

- obsolete element
 - in nativestore, 120, 122
- OpenJMS, 93
- OpenLDAP, 22, 51
 - disabling, 45
 - enabling, 45
- operating systems
 - troubleshooting, 115

- Oracle
 - errors, 116
- Oracle AQ, 93
- out of memory errors, 114
- overview, 12–13, 21

- pager , 64
- pages
 - configuration, 55–60, 64, 66, 69–71
 - Data Service, 57
 - Deployment Portal, 58
 - login, 12–13, 55
 - notification, 60
 - process management, 64
 - repository, 66
 - search, 69
 - ShowCase Warehouse Builder, 71
 - SMTP settings, 60
- password attribute
 - for user, 121
- passwords, 114
 - changing, 8, 13
 - providing, 12
 - supplying, 12
- PASW BIRT Report Designer, 3, 6
- PASW Modeler, 6
- PASW Modeler adapter, 115
- PASW Modeler adapter file permissions, 115
- PASW Statistics
 - credentials, 57
 - custom dialogs, 57
 - server, 57
- PDF, 115
- PEB report processing errors, 114
- performance bottlenecks, 113
- performance tuning, 87
- permanent generation heap size, 114
- persistent event queue, 88
- port numbers, 19
- process management
 - configuration, 64
- prompted credentials, 115
- protocol timeout, 57
- providers, 42

- query examples, 110
- queue, 93

- regulatory compliance, 97
- reindexing, 75
- reinstalling the repository, 114
- remote process
 - execution servers, 6
- removing
 - MIME types, 74
- reporting output, 115
- reports, 64
- repository, 3
 - configuration, 66
- repository events, 98
- repository servers
 - properties, 19
- rerunning setup, 114
- resolving errors, 113, 115–116
- RFC3461, 91
- role element
 - in user, 120–121
- roles, 22, 36
 - adding, 39
 - adding actions, 39
 - administrators, 38
 - assigning groups, 40
 - assigning users, 40
 - creating, 38
 - editing, 39
 - removing, 41
 - removing actions, 39

- sample batch file to run JMS client, 96
- sample JMS client program, 94
- sample JMS topic configuration file, 93
- SAP NetWeaver WebAS, 93
- SAS
 - execution server, 6
- schema
 - auditing database, 99
- scoring, 115
- scoring configuration, 58
- scoring service, 68, 115
- search, 69
- search limit, 69
- search service, 75
- security, 55, 69
- security events, 98
- security providers, 22, 42–43
 - Active Directory, 43, 49
 - Active Directory with local override, 44, 50
 - disabling, 49
 - enabling, 49
 - IBM i, 46
 - IBM i native, 50
 - IBM i user repository, 42
 - native, 43, 49
 - OpenLDAP, 45, 51
 - SiteMinder, 48, 51
- servers
 - starting, 8
 - stopping, 8, 11
 - UNIX, 8, 11
 - Windows, 8, 11
- services
 - UNIX, 8, 11
 - Windows, 8, 11

- session timeout, 69
- setup, 114
 - configuration, 70
- shared libraries, 114
- SHLIB_PATH, 114
- ShowCase Warehouse Builder
 - configuration, 71
- single sign-on, 42, 50, 52
- single sing-on, 12
- SiteMinder, 42, 48, 51
- siteminder.package, 42
- SMTP
 - logging, 91
 - message headers, 91
 - properties, 78
 - server threads, 88
- Solaris 10, 115
- Solaris 9
 - errors, 115
 - libm.so.1, 115
 - libm.so.2, 115
 - startup failure, 115
 - wrapper error, 115
- SQL queries, 97
- SSL, 19, 44, 50
- SSO, 12, 42, 50
- subscription identifiers cache, 88
- subscriptions management, 89
- Sun Java System Message Queue, 93
- SVG charts, 64
- symbolic link, 115
- system
 - configuring, 55–60, 64, 66, 68–71
 - launching, 8, 11–13
 - login, 8, 12–13
 - logout, 8
 - navigation, 8, 13
 - overview, 13, 21
 - starting, 8, 12–13
 - stopping, 11
- system errors, 113, 115–116
- system information, 14

- tablespaces, 116
- tabs
 - navigating, 13
- templates, 55–56
 - customizing content, 80
 - customizing format, 82
 - customizing properties, 78
 - for e-mail notifications, 77, 83
 - inserting event property variables, 80
 - inserting properties, 80
 - Velocity, 83
- timeout errors, 57
- topic, 93

- topics
 - naming, 21
- troubleshooting, 14, 113, 115–116
 - notification delivery failure, 91

- UNIX
 - services, 8, 11
- URL prefix, 70
- user element
 - in nativestore, 120
 - in obsolete, 122
- user preferences, 4
- userID attribute
 - for user, 121
- users
 - access to system resources, 22
 - adding, 23, 26
 - allowed, 22–23, 34
 - creating, 23, 26
 - deleting, 29
 - editing, 23, 28
 - group membership, 22
 - importing, 32
 - local, 22–23
 - managing in Deployment Manager, 22
 - modifying, 23, 28
 - remotely defined, 22–23
 - setting up, 22

- value-of element
 - in notification templates, 78, 80
- Velocity, 77
- version, 14
- viewing
 - job output, 87
 - server properties, 19
- visualization
 - reports, 66
 - specifications, 66

- WebSphere, 118
- WebSphere MQ, 93
- Windows
 - services, 8, 11
- wrapper error, 115
- wrapper service, 113