IBM SPSS Collaboration and Deployment Services
Version 6 Release 0

Administrator's Guide

Note
Before using this information and the product it supports, read the information in "Notices" on page 105.

Product Information

This edition applies to version 6, release 0, modification 0 of IBM SPSS Collaboration and Deployment Services and to all subsequent releases and modifications until otherwise indicated in new editions.

Contents

Chapter 1. Overview

IBM SPSS Collaboration and Deployment Services

IBM® SPSS® Collaboration and Deployment Services is an enterprise-level application that enables widespread use and deployment of predictive analytics.

IBM SPSS 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 users. The benefits of IBM SPSS 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

IBM SPSS 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 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, IBM SPSS Collaboration and Deployment Services provides a repository for storing assets using a folder hierarchy similar to most file systems. Files stored in the IBM SPSS Collaboration and Deployment Services 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 IBM SPSS Collaboration and Deployment Services. Products such as IBM SPSS Statistics and IBM SPSS 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 correct 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 IBM SPSS Collaboration and Deployment Services, individual files stored in the repository can be included in processing **jobs**. Jobs define an execution sequence for analytical artifacts and can be created with IBM SPSS Collaboration and Deployment Services Deployment Manager. 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 IBM SPSS Collaboration and Deployment Services 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 IBM SPSS 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 IBM SPSS 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, IBM SPSS Collaboration and Deployment Services consists of a single, centralized IBM SPSS Collaboration and Deployment Services Repository that serves a variety of clients, using execution servers to process analytical assets.

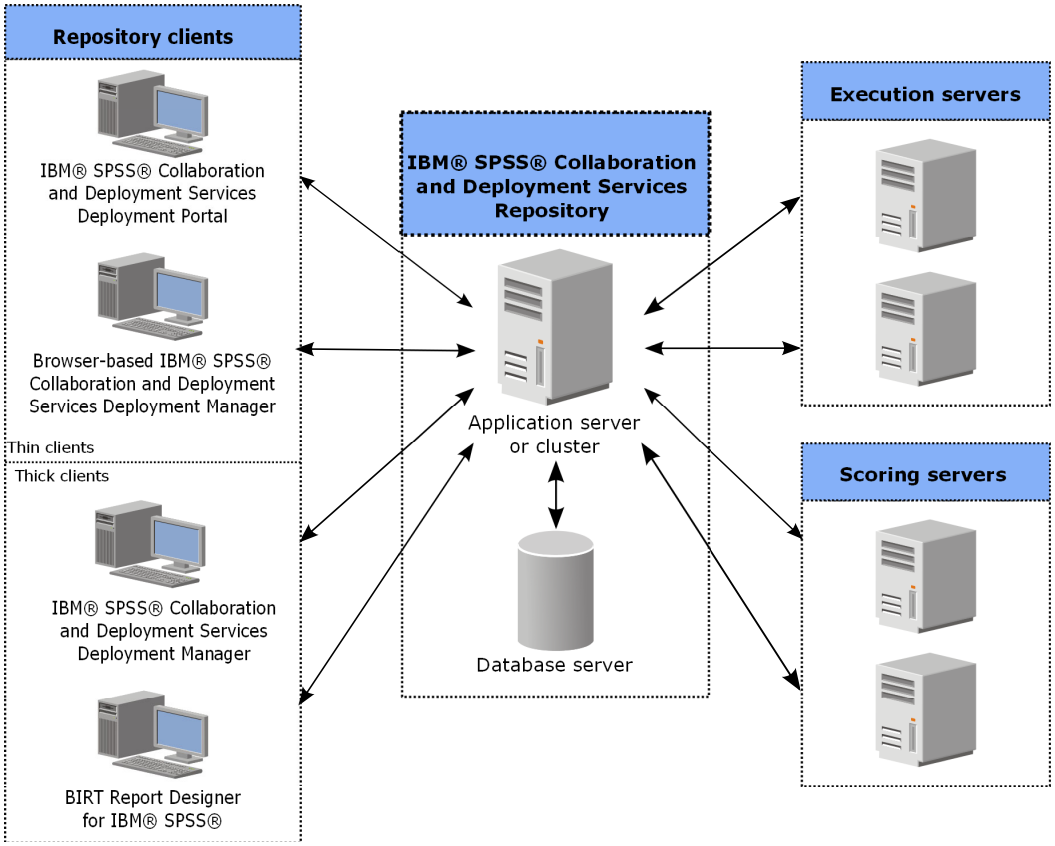

Figure 1. IBM SPSS Collaboration and Deployment Services Architecture

IBM SPSS Collaboration and Deployment Services consists of the following components:

- IBM SPSS Collaboration and Deployment Services Repository for analytical artifacts
- IBM SPSS Collaboration and Deployment Services Deployment Manager
- IBM SPSS Collaboration and Deployment Services Deployment Portal
- Browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager
- IBM SPSS Collaboration and Deployment Services Enterprise View
- BIRT Report Designer for IBM SPSS

IBM SPSS Collaboration and Deployment Services Repository

The repository provides a centralized location for storing analytical assets, such as models and data. The repository requires an installation of a relational database, such as IBM DB2, Microsoft SQL Server, or Oracle.

The repository includes facilities for:

- Security
- Version control
- Searching
- Auditing

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

IBM SPSS Collaboration and Deployment Services Deployment Manager

IBM SPSS Collaboration and Deployment Services Deployment Manager is a client application for IBM SPSS Collaboration and Deployment Services Repository that enables users to schedule, automate, and execute analytical tasks, such as updating models or generating scores.

The client application allows a user to perform the following tasks:

- 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
- Determine the status of a job
- Specify email notification of job status

In addition, the client application allows users to perform administrative tasks for IBM SPSS Collaboration and Deployment Services, including:

- Manage users
- Configure security providers
- Assign roles and actions

Browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager

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

- Setting system configuration options
- 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.

You typically access the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager at the following URL:

```
http://<host IP address>:<port>/security/login
```

Note: An IPv6 address must be enclosed in square brackets, such as `[3ffe:2a00:100:7031::1]`.

If your environment is configured to use a custom context path for server connections, include that path in the URL.

```
http://<host IP address>:<port>/<context path>/security/login
```

IBM SPSS Collaboration and Deployment Services Deployment Portal

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

The web 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 email address and password, general options, subscriptions, and options for output file formats

You typically access the home page at the following URL:

```
http://<host IP address>:<port>/peb
```

Note: An IPv6 address must be enclosed in square brackets, such as `[3ffe:2a00:100:7031::1]`.

If your environment is configured to use a custom context path for server connections, include that path in the URL.

```
http://<host IP address>:<port>/<context path>/peb
```

IBM SPSS Collaboration and Deployment Services Enterprise View

The IBM SPSS Collaboration and Deployment Services Enterprise View provides a single, consistent view of enterprise data. It 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, or an online transaction database. The Enterprise View ensures a consistent use of enterprise data and hides the complexities of stored data structures from the user. The 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. The 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.

The Enterprise View includes the IBM SPSS Collaboration and Deployment Services Enterprise View Driver component designed to provide other applications access to Enterprise View objects stored in the repository. The driver operates similarly to JDBC and ODBC drivers with the exception that it does not directly query physical data sources but rather virtualizes the physical data sources according to the design of the Data Provider Definitions. Note that while the Enterprise View is installed as part of IBM SPSS Collaboration and Deployment Services Deployment Manager, the IBM SPSS Collaboration and Deployment Services Enterprise View Driver driver must be installed separately. For more information, see IBM SPSS Collaboration and Deployment Services Enterprise View Driver documentation.

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 IBM SPSS 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.

Scoring server

IBM SPSS Collaboration and Deployment Services Scoring Service is also available as a separately deployable application, the Scoring Server.

The Scoring Server improves deployment flexibility in several key areas:

- Scoring performance can be scaled independently from other services
- Scoring Server(s) can be independently configured to dedicate computing resources to one or any number IBM SPSS Collaboration and Deployment Services scoring configurations
- Scoring Server operating system and processor architecture does not need match the IBM SPSS Collaboration and Deployment Services Repository or other Scoring Servers
- Scoring Server application server does not need match the application server used for IBM SPSS Collaboration and Deployment Services Repository or other Scoring Servers

BIRT Report Designer for IBM SPSS

The reporting functionality of IBM SPSS 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.

The IBM SPSS 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 IBM SPSS Collaboration and Deployment Services reporting job steps. BIRT Report Designer for IBM SPSS is a standalone application that can be used in conjunction with IBM SPSS 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 BIRT Report Designer for IBM SPSS report requires a JDBC-based database connection, a corresponding JDBC driver must be installed with the IBM SPSS Collaboration and Deployment Services Repository. For application server-specific information on the location of the JDBC drivers, see the corresponding section of the repository installation instructions.

To start BIRT Report Designer for IBM SPSS, run the file *BIRT.exe* in the installation directory. For information on using BIRT Report Designer for IBM SPSS, see the documentation installed with the application.

IBM Analytical Decision Management

IBM SPSS Collaboration and Deployment Services are a prerequisite for installing IBM Analytical Decision Management, a suite of applications for integrating predictive analytics with operational decision-making. IBM Analytical Decision Management uses high-speed scoring, master data management, and process automation facilities of IBM SPSS Collaboration and Deployment Services to optimize and automate high-volume decisions and produce improved outcomes in specific business situations.

Chapter 2. What is new in this release

What is new for administrators

IBM SPSS Collaboration and Deployment Services 6 delivers new capabilities that can help you simplify deployment of predictive analytics and manage costs.

IPv6 multicast address support

You can now refer to the IBM SPSS Collaboration and Deployment Services Repository server by using a multicast address of the IPv6 type.

Chapter 3. Getting started

After successfully installing the IBM SPSS Collaboration and Deployment Services 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 IBM SPSS Modeler support

Starting the repository server

The repository server 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, the repository server typically runs in the background, handling requests from clients such as IBM SPSS Modeler or the IBM SPSS Collaboration and Deployment Services 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.

Important: To avoid permissions conflicts, the repository server must always be started under the same credentials, preferably a user with `sudo` (UNIX) or administrator-level (Windows) privileges.

The repository server is started by starting the application server. This can be accomplished with the scripts provided with the repository server installation or native application server administration tools. For more information, see the application server vendor documentation.

WebSphere

Use WebSphere administration tools. For more information, see WebSphere documentation.

JBoss

Use the following scripts with the repository server installation:

```
<repository installation directory>/bin/startserver.bat
<repository installation directory>/bin/startserver.sh
```

Alternatively, you can also use JBoss administration tools to start the server. For more information, see JBoss documentation.

WebLogic

For single WebLogic server configurations, use the following scripts provided with the repository server installation:

```
<repository installation directory>/bin/startserver.bat
```

```
<repository installation directory>/bin/startserver.sh
```

The WebLogic application server can also be started using your preferred mechanism, but you must ensure the correct environment variables and Java properties are set. To assist with this process, the configuration process creates the following scripts in the *toDeploy/current* directory:

- *setCDSEnv.cmd* or *setCDSEnv.sh*
- *startCDSWebLogic.cmd* or *startCDSWebLogic.sh*
- *startManagedCDSWebLogic.cmd* or *startManagedCDSWebLogic.sh*

If you selected automatic deployment during the configuration, the files are also copied to the domain and *<domain>/bin* directory. Inspect these files to determine which environment and Java properties must be set. The specific properties will vary depending on installed IBM SPSS adapters. If you are starting your server using a startup script, you can call *setCDSEnv.cmd/setCDSEnv.sh* from that script. If you are using node manager or some other mechanism to start the server, make sure to define the equivalent settings.

Using the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager

The Login page is your gateway into the system.

To log in:

1. In a browser, navigate to the Login page. Typically, the URL is the following:
   ```
   http://<host IP address>:<port>/security/login
   ```
   The 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.

   Note: An IPv6 address must be enclosed in square brackets, such as `[3ffe:2a00:100:7031::1]`.

   If your environment is configured to use a custom context path for server connections, include that path in the URL.
   ```
   http://<host IP address>:<port>/<context path>/security/login
   ```
2. In the Login Name field, enter your user ID.
3. In the Password field, enter your password.
4. Click **Login**.

Important: To successfully log in, your browser must allow session cookies.

Additional options

On the Login page, you also have the option of changing your password. See the topic “Changing passwords” for more information.

Important: Browser-based IBM SPSS Collaboration and Deployment Services 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.

1. In the Login Name field, enter your login name.
2. In the Current Password field, enter your current password.

3. In the New Password field, enter your new password.
4. In the Confirm New Password field, reenter your new password.
5. Click **Save New Password**. In the Messages section, the following text appears:

   ```
   Password updated
   ```

6. Click **Return to Login**. The Login page opens. Log in to the system using your new password. See the topic "Using the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager" on page 10 for more information.

Navigating through the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager

The browser-based IBM SPSS Collaboration and Deployment Services 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, you can choose any of the following categories:

- **Configuration**
- **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 an item, its corresponding section appears in the right pane. If a section has multiple subsections, a series of tabs appears in the right 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 IBM SPSS 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.

To display detailed information for installed packages:

- Click **Show Details**.

To download a text file of version and system information:

- Click **Download version and system details**.

To download text files of version and system information and application server log:

- Click **Download version, system details and logs in one zip file**. The files are downloaded as compressed archive.

Using the IBM SPSS Collaboration and Deployment Services Deployment Manager

Administrative tasks can be performed using the IBM SPSS Collaboration and Deployment Services Deployment Manager as well as the browser-based IBM SPSS Collaboration and Deployment Services 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, IBM SPSS Collaboration and Deployment Services Deployment Manager allows administration of other servers, such as IBM SPSS Statistics and IBM SPSS Modeler servers.

Getting started

Administered servers

Server administration in IBM SPSS Collaboration and Deployment Services 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 IBM SPSS Collaboration and Deployment Services Deployment Manager sessions, facilitating access to those servers.

From the menus choose:

Tools > **Server Administration**

The administered server list may include a variety of server types, including IBM SPSS Collaboration and Deployment Services Repository servers, IBM SPSS Modeler servers, and IBM SPSS 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 IBM SPSS 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 the name and type for the server.

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
- Forward slash (/)
- Periods
- Commas
- Semicolons

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

- IBM SPSS Collaboration and Deployment Services Repository Server
- Administered IBM SPSS Modeler Server
- Administered IBM SPSS Statistics Server
- Administered IBM SPSS Modeler 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 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.

For an IBM SPSS Collaboration and Deployment Services Repository server, you can specify the server URL.

The URL includes the following elements:

- The connection scheme, or protocol, as either *http* for hypertext transfer protocol or *https* for hypertext transfer protocol with the secure socket layer (SSL)
- The host server name or IP address

  Note: An IPv6 address must be enclosed in square brackets, such as `[3ffe:2a00:100:7031::1]`.
- The port number. If the repository server is using the default port (port 80 for http or port 443 for https), the port number is optional.
- An optional custom context path for the repository server

Table 1. Example URL specifications. This table lists some example URL specifications for server connections.

URL	Scheme	Host	Port	Custom path
`http://myserver`	HTTP	*myserver*	default (80)	(none)
`https://9.30.86.11:443/spss`	HTTPS	9.30.86.11	443	*spss*
`http://[3ffe:2a00:100:7031::1]:9080/ibm/cds`	HTTP	`3ffe:2a00:100:7031::1`	9080	*ibm/cds*

Contact your system administrator if you are unsure of the URL to use for your server.

For other server types, available properties include the following items:

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
- Forward slash (/)
- Periods
- Commas
- Semicolons

Port The port number that is 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.

After you define the properties, the new server is included 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 displayed properties depend on the type of server selected.

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.

IBM SPSS Collaboration and Deployment Services Repository Server Login

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 your administrative tasks, log off from the server.

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

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.

Chapter 4. Users and groups

An IBM SPSS 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. See the topic "Roles overview" on page 25 for more information.

IBM SPSS 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, IBM SPSS Collaboration and Deployment Services itself is the security provider) or derived from a remote directory, such as Windows Active Directory or OpenLDAP. See the topic Chapter 7, "Security providers," on page 31 for more information.

Some environments may require setting up groups of remotely defined users that are specific to IBM SPSS Collaboration and Deployment Services 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 IBM SPSS 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 IBM SPSS Collaboration and Deployment Services users

Local user setup in IBM SPSS Collaboration and Deployment Services involves:

1. Creating the user and, if necessary, assigning group membership. Local user and groups can be managed through IBM SPSS Collaboration and Deployment Services Deployment Manager.
2. Defining the user's level of access by assigning the role on an individual or group basis. See the topic "Editing users and groups assigned to a role" on page 27 for more information. If the role with an appropriate set actions does not exist, it must be established. See the topic "Creating a new role" on page 27 for more information.

Externally defined user setup in IBM SPSS 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. See the topic "Configuring security providers" on page 32 for more information.

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 IBM SPSS Collaboration and Deployment Services 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 IBM SPSS Collaboration and Deployment Services 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 IBM SPSS Collaboration and Deployment Services Deployment Manager

IBM SPSS Collaboration and Deployment Services 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.

Before 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 IBM SPSS Collaboration and Deployment Services 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.
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, import users and groups, and lock or unlock users from accessing the system.
- For Active Directory with Local Override, the editor displays all externally defined groups and users that have been set up to access IBM SPSS 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. See the topic Chapter 7, “Security providers,” on page 31 for more information.

Creating a user

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

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.

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.

Locking and unlocking a user

By default, after a native Local User Repository user attempts to log on to IBM SPSS Collaboration and Deployment Services with an incorrect password three consecutive times, their user account will be locked automatically. The user will be unable to log in (even with the correct credentials) until their account is unlocked automatically after thirty minutes or unlocked manually by an administrator.

In the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager, under the Security section, there are two configuration settings to customize this functionality:

- **Invalid Login Attempt Count Threshold.** This setting defines the number of times to allow a failed login before automatically locking out the user. You can also choose to never lock users automatically.
- **Account Lockout Duration.** This setting defines the number of minutes to wait before automatically unlocking users who have been locked out. You can also choose to never unlock users automatically.

Note that this functionality only applies to Local User Repository native security provider users.

In the Manage Users and Groups editor for Local User Repository, you can also lock and unlock local users manually. The State column indicates if a user is locked. To display all users who are currently locked, select **Show only locked users** in the Manage Users and Groups editor.

To manually unlock a local user:

1. Select the locked user in the Manage Users and Groups editor. The State column displays the text **Locked** for any users who are locked. If you want to view all users currently locked, click **Show only locked users**.
2. Click **Unlock**. A dialog box opens to confirm that the user should be unlocked.
3. Click **Yes** to unlock the user.

To manually lock a local user:

1. Select the user you want to lock in the Manage Users and Groups editor. You cannot lock groups.
2. Click **Lock**. A dialog box opens to confirm that the user should be locked.
3. Click **Yes** to lock the user. Note that a user who is manually locked will remain locked out until unlocked manually. The Account Lockout Duration configuration setting described previously is not applied (the user will not be unlocked automatically).

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.

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**.

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 Chapter 16, “nativestore schema reference,” on page 101.

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.
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.

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.

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.

Chapter 5. 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, whereas roles are managed across security providers.

Use the Server Administration tool of IBM SPSS Collaboration and Deployment Services 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.

IBM SPSS Collaboration and Deployment Services actions

- **Access Contents and Folders.** Access the IBM SPSS Collaboration and Deployment Services Repository.
- **Access Syndicated Feeds.** Access syndicated feeds such as RSS (Really Simple Syndication) feeds.
- **Configuration.** Modify repository settings.
- **Configure Model.** Configure models for scoring.
- **Create Subscriptions.** Create individual subscribtions 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 the repository.
- **Define Datasources.** Define and modify data sources.
- **Define Message Domains.** Define and modify domains for JMS messaging.
- **Define Promotion Policies.** Define and modify policies (sets of rules) for promoting repository objects.
- **Define Server Clusters.** Define and modify execution server clusters.
- **Define Servers.** Define and modify execution servers.
- **Define Topics.** Define and modify topic hierarchy for the 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 IBM SPSS Collaboration and Deployment Services 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 the repository.
- **Promote Objects.** Promote repository objects.
- **Repository Index.** Reindex the contents of the repository.
- **Run Custom Dialogs** Run IBM SPSS Statistics custom dialogs.
- **Run Report Dynamically.** Run dynamic reports, such as Business Intelligence Reporting Tools (BIRT) reports, in IBM SPSS Collaboration and Deployment Services Deployment Portal.
- **Schedules.** Manage job schedules.
- **Score Model.** Score models.
- **Show All Versions.** View all versions of objects (labeled and unlabeled) in IBM SPSS Collaboration and Deployment Services Deployment Portal. By default, users are able to see only labeled versions in IBM SPSS Collaboration and Deployment Services Deployment Portal.
- **Show latest.** View only the latest version of objects.
- **Submit Work** Submit work (for example, reports) for processing by IBM SPSS Collaboration and Deployment Services.
- **User Preference Administration.** Manage the preferences of other users. Note that IBM SPSS Collaboration and Deployment Services products do not provide any user interfaces for modifying the preferences of other users. This setting only applies if calling the User Preference Web Service directly.
- **View Expired Files.** View expired content, such as files and jobs.
- **View Model Management Dashboard.** View model management dashboards in IBM SPSS Collaboration and Deployment Services Deployment Manager and IBM SPSS Collaboration and Deployment Services 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. See the topic “Creating a new role” on page 27 for more information.

Manage role definitions

To work with roles, choose **Server Administration** from the **Tools** menu, select an IBM SPSS Collaboration and Deployment Services Repository Server, and log in. Double-click the **Roles** icon for the server to access the Manage Role Definitions editor.

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 role, click the **New Role** button in the Roles editor. A role needs a name and a list of associated actions.

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.

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

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.

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.

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

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.

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

The role is removed from the system.

Chapter 6. Cross Site Scripting (XSS) filters

Cross Site Scripting (XSS) is a computer security vulnerability typically found in web applications. It enables attackers to bypass client-side security mechanisms normally imposed on web content by modern web browsers by injecting malicious script into web pages viewed by other users.

XSS can be a significant security risk depending on the sensitivity of your data. In versions of IBM SPSS Collaboration and Deployment Services previous to 5.0.0.0, a web security filter was available to help prevent XSS attacks by validating user-entered parameters. But all the filter criteria were embedded in the product and not available for editing or customization by users. With IBM SPSS Collaboration and Deployment Services Deployment Manager, users can now add, modify, and delete XSS filter rules based on their company's enterprise security policy.

Managing XSS filter rules

IBM SPSS Collaboration and Deployment Services Deployment Manager allows you to manage XSS filter rules based on your company's enterprise security policy. To work with XSS filters, first navigate to the administrative interface:

1. From the **Tools** menu, choose **Server Administration**.
2. On the Server Administration tab, log in to a repository server. Double-click the **Configuration** icon to expand the hierarchy.
3. Double-click the **Cross Site Scripting Filters** icon.

The Manage XSS Filter Rule Definitions editor opens.

The editor displays all XSS filter rules currently defined for the server. Administrators can create, modify, and delete XSS filter rules. Select a filter type from the drop-down to display any filter rules currently defined for that type. The following filter types are available.

- Restrict HTML Elements
- Restrict JavaScript functions
- Restrict plain text strings
- Regular expressions for restrict string
- Allowed strings

Changes to XSS filter rules are applied immediately (restarting the server is not required).

Creating XSS filter rules

To create a new XSS filter rule:

1. In the Manage XSS Filter Rule Definitions editor, select the filter type for which you want to create a new rule.
2. Click **Add**. The Edit Rule dialog opens.
3. Type the value for the new XSS filter rule and click **OK**.

This documentation does not provide any example XSS filter rules. Doing so might provide ideas for malicious scripts.

Chapter 7. Security providers

A security provider is responsible for verifying the credentials supplied by a user against a particular user directory. IBM SPSS 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 IBM SPSS Collaboration and Deployment Services, in which users, groups, and roles can all be defined. The native provider is always active and cannot be disabled.
- **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 IBM SPSS 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. In contrast to the native provider, this provider can be enabled or 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 IBM SPSS 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. 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 IBM SPSS Collaboration and Deployment Services users. This provider can be enabled or disabled. If IBM i security provider is used with single sign-on-enabled IBM SPSS Collaboration and Deployment Services installation, EIM (Enterprise Identity Management) must be enabled. Additionally, */QIBM/UserData/Java400/ext/eim.jar* must be copied into the library directory of the IBM SPSS Collaboration and Deployment Services application server if the application server is running on a non-IBM i host.

Security providers in IBM SPSS Collaboration and Deployment Services 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 an IBM SPSS 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. To configure a new security provider, right-click **Security Providers** and select **New**. A wizard will be displayed. Or to modify an existing security provider configuration, double-click the security provider name under **Security Providers**.

To enable or disable security providers, right-click them on the Server Administration tab and select **Enable** or **Disable**.

Configuring security providers

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

See the following topics for details.

- Native
- OpenLDAP
- Active Directory
- Active Directory with local override
- IBM i

To enable or disable security providers, right-click them on the Server Administration tab and select **Enable** or **Disable**.

Note: When changes are made to an already existing security provider definition, they are not activated until the repository is restarted or until the security provider is disabled and re-enabled. In certain cases, for example when the domain name for Active Directory security provider is changed, users and groups must be removed and re-added to roles. See the topic "Setting up IBM SPSS Collaboration and Deployment Services users" on page 17 for more information.

Native

The Local User Repository native security provider is internal to IBM SPSS Collaboration and Deployment Services and does not contain any settings to configure.

OpenLDAP

To modify an existing OpenLDAP configuration, double-click the **OpenLDAP** entry under **Security Providers**.

To configure a new OpenLDAP security provider, right-click **Security Providers** and select

New > **Security Provider Definition**

The Create New Security Provider Definition wizard will be displayed. Select **OpenLDAP** from the **Type** drop-down menu. Type a name for the security provider definition, click **Next**, and proceed through the steps in the wizard. Refer to the following details.

Host settings

- **Host URL.** The path to the LDAP server, usually a DNS resolvable name or an IP address (for example, *ldap://yourserver.yourcompany.com*). The default port for LDAP is 389.
- **Use Secured Socket Layer connection.** Select to use secure sockets for communication with the OpenLDAP server.
- **Page Search Result.** Select this option if your LDAP server provides an option for paging LDAP search output, and only when this option is enabled. Additional information about the paged results search control can be found in *RFC 2686 - LDAP Control Extension for Simple Paged Results Manipulation* (*http://datatracker.ietf.org/doc/rfc2696/*).

Credentials

- **Search Credential Type.** Specify the handling of search credentials. When the back end server allows it, the *Use Anonymous Bind* option provides the ability to search for users without having to provide a search user ID and search user password. The *Use Kerberos Credential* option uses the system Server Process Credential for searches. Select the *Use Supplied Credential* option to specify a user identifier and password to use as search credentials.

- **Search user.** 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.
- **Search 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.

User bind definition

- **Use Context Bind.** Select to perform a bind operation when the user logs in. This is recommended.
- **Password Attribute.** The password attribute to use when User Bind is not desired. If chosen, you are confirming that the security server allows a return value of the password attribute in queries. Otherwise this option cannot be used.
- **Password Digest.** The password digest method used by the security server to hash password. This option is used if User Bind is not desired. If chosen, you are confirming that the security server allows a return value of the password attribute in queries. Otherwise this option cannot be used.

User search settings

- **Search Filter Base DN.** Base distinguished name for user searches.
- **Object Filter Expression.** The object class and value to use for filtering. This value is dependent on the LDAP schema being used.
- **Search Filter Expression.** The attribute to use as the search ID. This value is dependent on the LDAP schema being used.
- **Search on Attribute.** The attribute that matches the Search Filter Expression attribute. This value is dependent on the LDAP schema being used.
- **Group User Filter.** Attribute indicating user group membership.

Group search settings

- **Search Filter Base DN.** Base distinguished name for group searches.
- **Object Filter Expression.** The object class and value to use for filtering. This value is dependent on the LDAP schema being used.
- **Search Filter Expression.** The attribute to use as the search ID. This value is dependent on the LDAP schema being used.
- **Group Attribute.** The attribute that matches Search Filter Expression attribute. This value is dependent on the LDAP schema being used.
- **Membership Attribute.** The attribute that denotes group membership. This value is dependent on the LDAP schema being used.
- **Refresh Interval.** Interval at which group membership data is refreshed.

Active Directory

To configure a new Active Directory security provider, right-click **Security Providers** and select

New > Security Provider Definition

The Create New Security Provider Definition wizard will be displayed. Select **Active Directory** from the **Type** drop-down menu. Type a name for the security provider definition, click **Next**, and proceed through the steps in the wizard. Refer to the following details.

Host settings

- **Host URL.** URL for the Active Directory server. The default port for LDAP is 389.
- **Use Secured Socket Layer connection.** Select to use secure sockets for communication with the Active Directory server.
- **Page Search Result.** Select this option if your Active Directory server provides an option for paging Active Directory search output, and only when this option is enabled.

Credentials

- **Search Credential Type.** Specify the handling of search credentials. When the back end server allows it, the *Use Anonymous Bind* option provides the ability to search for users without having to provide a search user ID and search user password. The *Use Kerberos Credential* option uses the system Server Process Credential for searches. Select the *Use Supplied Credential* option to specify a user identifier and password to use as search credentials.
- **Search 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.
- **Search 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.

Domain name

- **Domain.** The DNS namespace to which the user is logging in.

User bind definition

- **Use Context Bind.** Select to perform a bind operation when the user logs in. This is recommended.
- **Password Attribute.** The password attribute to use when User Bind is not desired. If chosen, you are confirming that the security server allows a return value of the password attribute in queries. Otherwise this option cannot be used.
- **Password Digest.** The password digest method used by the security server to hash password. This option is used if User Bind is not desired. If chosen, you are confirming that the security server allows a return value of the password attribute in queries. Otherwise this option cannot be used.

User search settings

- **Search Filter Base DN.** Base distinguished name for user searches.
- **Object Filter Expression.** The object class and value to use for filtering. This value is dependent on the schema being used.
- **Search Filter Expression.** The attribute to use for the search ID. This value is dependent on the schema being used.
- **Search on Attribute.** The attribute that matches the Search Filter Expression attribute. This value is dependent on the schema being used.
- **Group User Filter.** Attribute indicating user group membership.

Group search settings

- **Search Filter Base DN.** Base distinguished name for group searches.
- **Object Filter Expression.** The object class and value to use for filtering. This value is dependent on the LDAP schema being used.
- **Search Filter Expression.** The attribute to use as the search ID. This value is dependent on the LDAP schema being used.
- **Group Attribute.** The attribute that matches Search Filter Expression attribute. This value is dependent on the LDAP schema being used.
- **Membership Attribute.** The attribute that denotes group membership. This value is dependent on the LDAP schema being used.
- **Refresh Interval.** Interval at which group membership data is refreshed.

Active Directory with local override

To configure a new Active Directory with local override security provider, right-click **Security Providers** and select

New > Security Provider Definition

The Create New Security Provider Definition wizard will be displayed. Select **Active Directory with Local Override** from the **Type** drop-down menu. Type a name for the security provider definition, click **Next**, and proceed through the steps in the wizard.

Most of the settings are identical to the Active Directory settings. However, local override offers two additional settings:

- **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.

IBM i

After installation, IBM i will be displayed under Security Providers in IBM SPSS Collaboration and Deployment Services Deployment Manager. To configure the IBM i security provider, specify values for the following settings:

- **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 IBM SPSS 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 IBM SPSS 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 of 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 profile's 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 IBM SPSS 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 IBM SPSS 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 IBM SPSS Collaboration and Deployment Services. However, if *test2* is associated with *test* before IBM SPSS Collaboration and Deployment Services permissions of *test* are defined, *test2* does inherit the permissions.

Security providers in the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager

To access the Security Providers page:

1. Click **Security Providers** in the navigation list. The Security Providers page appears.

   To modify the security providers used:

2. Select (to enable) or clear (to disable) the check boxes next to the security provider.
3. Click **Set**.

Note that only security providers that have first been created in IBM SPSS Collaboration and Deployment Services Deployment Manager client will appear in the list.

Enabling and disabling security providers

Only security providers that have first been created and configured in IBM SPSS Collaboration and Deployment Services Deployment Manager client will appear in the browser. For each type of security provider, you can view some settings specific to the type of authentication and authorization system being used. But to configure new security providers or modify the full set of settings, use the IBM SPSS Collaboration and Deployment Services Deployment Manager client.

You can enable or disable the available security providers by using the check boxes next to each security provider and clicking **Set**.

Native (local)

The native (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.

Active Directory

To view certain Active Directory settings, click **View settings** to the right of the Active Directory check box. A subset of the current settings appear.

Note that the Active Directory security provider is only available if it has first been configured in IBM SPSS Collaboration and Deployment Services Deployment Manager client. For information about specific settings, see "Active Directory" on page 33.

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 view certain Active Directory with Local Override settings, click **View settings** to the right of the Active Directory with Local Override check box. A subset of the current settings appear. Most of the settings correspond to those for Active Directory. However, the following two options are also available. Note that the Active Directory with local override security provider is only available if it has first been configured in IBM SPSS Collaboration and Deployment Services Deployment Manager client.

- **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.

IBM i

When the IBM SPSS Collaboration and Deployment Services Repository is installed on IBM i, the IBM i user profile directory will be used to authenticate repository logins.

To view certain IBM i security settings, click **View settings** to the right of the IBM i check box. A subset of the current settings appear. Note that the IBM i security provider is only available if it has first been installed and configured in IBM SPSS Collaboration and Deployment Services Deployment Manager client.

- **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 IBM SPSS 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.
- **Enable EIM Lookup.** For single sign-on enabled IBM SPSS 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 of 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 view certain OpenLDAP settings, click **View settings** to the right of the OpenLDAP check box. A subset of the current settings appear. Note that the OpenLDAP security provider is only available if it has first been configured in IBM SPSS Collaboration and Deployment Services Deployment Manager client. For information about specific settings, see “OpenLDAP” on page 32.

Chapter 8. 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.

IBM SPSS Collaboration and Deployment Services provides single sign-on capability by initially authenticating users through an external directory service based on the *Kerberos* security protocol, and subsequently using the credentials in all IBM SPSS Collaboration and Deployment Services applications (for example, IBM SPSS Collaboration and Deployment Services Deployment Manager, IBM SPSS Collaboration and Deployment Services Deployment Portal, or a portal server) without additional authentication.

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

Configuring single sign-on

1. Choose **Server Administration** from the **Tools** menu, log in to a IBM SPSS Collaboration and Deployment Services server, and double-click the **Single Sign-On** icon. Single Sign-on Provider editor opens.

- **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.
- **Host.** The name of the IBM SPSS Collaboration and Deployment Services Repository host. For example, `repositoryhost.mycompany.com`.
- **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 IBM SPSS 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.

Chapter 9. Repository configuration

IBM SPSS 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 IBM SPSS Collaboration and Deployment Services Deployment Manager:

1. Click **Configuration** in the navigation list. The Configuration page opens.
2. 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 are intended for optional IBM SPSS Collaboration and Deployment Services components or other IBM SPSS products, such as IBM SPSS Statistics. 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:

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

BIRT Report Designer for IBM SPSS

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 BIRT Report Designer for IBM SPSS in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 2. BIRT Report Designer for IBM SPSS configuration options.

Name	Description	Settings
BIRT Linked Resource Location	The directory on the server file system in which external resources for reports, like cascading style sheets 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.

Table 2. BIRT Report Designer for IBM SPSS configuration options (continued).

Name	Description	Settings
Enable SVG Chart	Specifies whether SVG chart output should be enabled. This setting should only be selected when SVG output is needed, and the browsers viewing the report output are SVG enabled. If disabled, reports that would generate charts use the PNG image format instead of SVG.	Disabled by default.

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 IBM SPSS 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 the IBM SPSS Collaboration and Deployment Services Repository cache. To enable Oracle Coherence:

1. Obtain and license Coherence components from Oracle. Coherence JAR files and all prerequisites must be placed into the *<IBM SPSS Collaboration and Deployment Services installation location>/components/cache-provider*.
2. Install the *coherence_cache_provider.package* from the *optional* folder on the IBM SPSS Collaboration and Deployment Services installation disc.
3. 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. **Note:** For IPv6-enabled local networks, use an address in IPv6 format. For example, you can specify `ff02::1` to multicast to all nodes on the local network segment.
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 *Multicast Group Address* and *Multicast Group Port* values to override the defaults.	Disabled by default.

Coordinator of Processes

Coordinator of Processes configuration options allow you to specify settings affecting the expiration time limit for connection requests and maintenance activities for the Coordinator of Processes.

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

Table 3. Coordinator of Processes configuration options.

Name	Description	Settings
Pending Connection Timeout	The expiration time limit for pending connection requests. The Coordinator of Processes will discard a connection request if the targeted server does not respond within the specified time interval.	Integer value. Default is 5 (seconds).
Coordinator of Processes Maintenance Provider Enabled	Enables or disables maintenance activities for the Coordinator of Processes	Enabled by default.

Custom dialog

If available, the IBM SPSS Statistics custom dialog configuration options allow you to specify settings for running custom dialogs.

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

Table 4. Custom Dialog configuration options.

Name	Description	Settings
File server browse enabled	Defines whether browsing for IBM SPSS Statistics data sets on the specified file server is enabled when selecting a data set for a custom dialog.	Check to enable.
File server location	The location of a file server (external to the repository) used to browse for IBM SPSS Statistics data sets when selecting a data set for a custom dialog. If file server browsing is enabled and no location is specified, then the file system of the specified IBM SPSS Statistics server will be used.	The value may be a network path or the absolute path of a directory.
File server name	A name to associate with the file server used to browse for IBM SPSS Statistics data sets.	A string value. If no value is specified then the name "File Server" is used.
Repository browse enabled	Defines whether browsing for IBM SPSS Statistics data sets in the repository is enabled when selecting a data set for a custom dialog.	Enabled by default.

Table 4. Custom Dialog configuration options (continued).

Name	Description	Settings
IBM SPSS Statistics server	The repository name or URI of an IBM SPSS 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 `spsscr:///?id=0a30063bc975ede400`. The URI can be found in the object properties. For more information, see IBM SPSS Collaboration and Deployment Services Deployment Manager documentation.
IBM SPSS Statistics server credential	The credential used to connect to the IBM SPSS Statistics server when executing custom dialog syntax. *Note*: The credential is not needed if Active Directory has been configured for use with IBM SPSS Collaboration and Deployment Services.	A string value corresponding to the repository name or URI of the credential object.
IBM SPSS Statistics server session timeout	Defines the timeout value, in minutes, for maintaining a connection to the IBM SPSS Statistics server in the absence of activity from a user.	Integer value. Default is 20 (minutes).

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 Data Service in the Configuration list. See the following table for link names, descriptions, and valid settings.

Table 5. 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 IBM SPSS Collaboration and Deployment Services Deployment Manager and the repository.

Specify the number of seconds the IBM SPSS Collaboration and Deployment Services Deployment Manager client should wait for a repository server. Use a larger value if timeout errors are received for server transactions.

To modify the protocol timeout:

1. In the Configuration list, under Deployment Manager, click **Protocol Timeout**. The current value appears.
2. In the Protocol Timeout text box, enter the number of seconds.
3. Click **Set**. The value you specified becomes the timeout value.
4. 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 for the web-based IBM SPSS Collaboration and Deployment Services 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 6. IBM SPSS Collaboration and Deployment Services Deployment Portal configuration options.

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

1. In the Configuration list, under Deployment Portal Scoring, click **Batch Scoring Row Limit**. The current value appears.
2. In the Batch Scoring Row Limit text box, enter the number of rows.
3. Click **Set**. The value you specified becomes the row limit.
4. 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 an IBM SPSS 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. IBM SPSS Collaboration and Deployment Services 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.
IBM SPSS Statistics data file Additional Servers	This setting is used to specify additional IBM SPSS Statistics data file servers that can be used to retrieve metadata from IBM SPSS Statistics data files.	A semicolon delimited list of `host:port` values, for example, `server2:18886;server3:18886`
IBM SPSS Statistics data file Load Balance	The load balancing setting controls whether multiple IBM SPSS Statistics data file servers are used in failover mode or load balancing mode when retrieving metadata from IBM SPSS 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 IBM SPSS Statistics data file servers are specified.	Enabled by default.
IBM SPSS Statistics data file Server Host	The name of the IBM SPSS Statistics data file server used for accessing IBM SPSS Statistics data files. If a host is not specified, the localhost will be used.	Any valid IP address or hostname.
IBM SPSS Statistics data file Server Port	The port for the IBM SPSS Statistics data file server. It the port is not specified, the default port will be used.	A valid port number.
IBM SPSS Statistics data file Server Secure	Indicator of whether or not SSL should be used when communicating with the IBM SPSS 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 options allow you to specify the location of the documentation components for browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager.

By default, the system uses paths established by the installation program. The Table 8 table describes the available settings.

Table 8. Help configuration options.

Name	Description	Settings
Guide Directory	Specifies the location of IBM SPSS Collaboration and Deployment Services guides and manuals.	The path of the directory that contains the guides.

Table 8. Help configuration options (continued).

Name	Description	Settings
Help Directory	Defines the location of the help system for IBM SPSS Collaboration and Deployment Services Deployment Manager.	The path of the directory that contains the help system.

To modify a help setting, perform the following steps:

1. In the Configuration list, click the setting to change from the **Help** group. The current value is shown.
2. Enter the new value.
3. Click **Set**. The value that you specified becomes the current value for that setting.

To return to the system-defined default, click **Use Default**. This option restores the default value 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.

See the topic "Optimizing notification service performance" on page 77 for more information. You can also specify syndication settings for feeds such as RSS (Really Simple Syndication).

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 9. Notification configuration options.

Name	Description	Settings
Binary Content Enabled	Enables binary content, such as email attachments, for notification messages.	Enabled 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.	Enabled by default.
Event Collector Enabled	Defines whether notification events should be processed by the service.	Enabled by default.
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.

Table 9. Notification configuration options (continued).

Name	Description	Settings
Event Inheritance Enabled	Defines whether derived notification events should be processed by the service.	Disabled 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.	Disabled 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.	Enabled by default.
Message Bus Filter Enabled	Defines whether only the notifications of interest should be sent to the JMS Message Bus.	Enabled by default.
Notification Auditor Enabled	Defines whether the notification service should interface with the auditing service.	Enabled by default.
Notification Cache Distributed	Defines whether the notification service should use a distributed cache. The notification service should be restarted for the changes to take effect.	Disabled 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.	Disabled 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.
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`.

Table 9. Notification configuration options (continued).

Name	Description	Settings
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 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 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. Clearing the check box will reverse the order of processing.	Enabled 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. The repository administrator can disable SMTP distribution to suppress all the emails generated by the server. Note that since the repository does not store generated email messages, if the SMTP distribution is disabled, all messages will be lost.	Enabled 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`.
SMTP EHLO	If false, do not attempt to sign on with the `EHLO` command.	True or false. Default is true.
SMTP from email address	The sender or return address to use for notification email.	Any existing SMTP email address.
SMTP Host	The IP address or hostname of the SMTP server used to send mail.	Any valid IP address or hostname.

Table 9. Notification configuration options (continued).

Name	Description	Settings
SMTP Local Host	Local hostname used in the SMTP `HELO` or `EHLO` command. Defaults to *InetAddress.getLocalHost().getHostName()*. 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 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 *SendFailedException*. 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 `smtp` or `smtps`. Default is `smtp` while`smtps` 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 Identifiers Cache	Defines a maximum size of the LRU cache for commonly used subscription identifiers.	Integer value. Default is 2048.
Syndicated Entry Cache TTL	Defines how long the syndicated feed entries will be saved in the cache (in minutes). This is for feeds such as RSS.	Integer value. Default is 15 minutes.
Syndicated Entry Max	Defines the maximum number of entries in the syndicated feeds, such as RSS.	Integer value. Default is 256.
Syndicated Entry Persistent TTL	Defines how long the syndicated entries will be saved in the persistent storage (in days). This is for feeds such as RSS.	Integer value. Default is 7 days.
Syndicated Feed Type	Defines the format for the syndicated feeds.	Either `RSS_2_0` or `ATOM_1_0`. Default is `RSS_2_0`.

Table 9. Notification configuration options (continued).

Name	Description	Settings
Syndication Distributor Enabled	Enables the syndication distributor for XML feeds.	Enabled by default.
Syndication Vacuumer Enabled	Enables the syndication vacuumer. The syndication vacuumer deletes expired syndicated entries from the system. It operates automatically based on the intervals specified in the Syndication Vacuumer Frequency option and uses the Syndicated Entry Persistent TTL value to determine what data is expired and available for deletion. Lack of vacuuming can seriously degrade application performance. Disabling this option is not recommended.	Enabled by default.
Syndication Vacuumer Frequency	Defines the frequency (in minutes) the syndication vacuumer will run. You must restart the Notification Service for changes to take effect.	Integer value. Default is 60 minutes.
Syndication Vacuumer Master	Defines whether the syndication vacuumer runs only on the master node in the server cluster.	Disabled by default.
Syndication Vacuumer Quota	Limits the number of syndicated entries to delete during a single run of the syndication vacuumer.	Integer value. Default is 4096.
URL data source 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.

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.

Important: You must restart the repository for the new option value to take effect.

To modify the pager timeout:

1. In the Configuration list, under Pager click **Pager Timeout**. The current value appears.
2. In the Pager Timeout text box, enter the number of minutes.
3. Click **Set**. The value you specified becomes the timeout value.
4. 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 for process management.

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 10. 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.
Hash-bang shell path	Specifies the hash-bang (#!) combination for the first line of the Unix script, followed by a pathname of the shell that executes the script.	
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 `ConnectionFactory`. 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 *org.jnp.interfaces.NamingContextFactory.* 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 `jnp://localhost:1099`. 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.
Job History Limit	Maximum number of job history entries to save for each version of a job. When the limit is reached, the oldest job history entries are replaced with new entries.	Integer value. Default is 10.
Log Query Metrics	Indicates whether or not to log query metrics (run time) to the log.	Disabled by default.
Maximum Number of Iterations	Maximum number of iterations for iterative job steps.	Integer value. Default is 250.
Message Poll	The length of time (in seconds) the Process Management server waits before repeating a scan of the repository for message-based schedules that should be activated.	Integer value. Default is 120.
Modeler Sync	Defines whether concurrent execution of jobs containing IBM SPSS Modeler files is allowed.	Disabled by default.

Table 10. Process management configuration options (continued).

Name	Description	Settings
Process Notification Enabled	Indicates whether the Process Management server should interface with the Notification Server.	True or false. Default is true.
Remote Process Server Poll	The length of time (in seconds) that remote work will wait before checking to see if the Remote Process Server is still active	
Remove Expired Submitted Artifacts	Indicates whether the artifacts created by submitting a resource for processing should be removed from the repository when they expire.	Enabled by default.
Remove Obsolete Job Histories	Indicates whether obsolete job histories should be removed.	Enabled by default.
Submitted Artifact Expiration Time	The expiration period (in days) for submitted artifacts such as jobs.	Integer value. Default is 5.
Submitted Artifact Timestamp	Timestamp format to be used in the names of Submitted Work folders generated automatically.	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:

1. In the Configuration list, under Reporting click **Complete Visualization Directory**. The current directory appears in the Complete Visualization Directory text box.
2. Enter the new value of the absolute path of the directory.
3. 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 11. Repository configuration options.

Name	Description	Settings
Categorical Value Limit	Limits the number of categorical variable values that are saved as IBM SPSS Modeler stream metadata. The saved values are included in the content evaluated when performing searches. The limit is necessary to decrease the time it takes to save a stream to the repository and perform searches.	Integer value. A value of -1 corresponds to no limit; all categorical values are saved as metadata. 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 repository.	URL.
Credential passwords must be encrypted	Credentials passwords must be encrypted. False indicates that passwords can be passed as unencrypted text. **Note:** This option is redundant for IBM SPSS Collaboration and Deployment Services deployments where SSL is already enabled and should be used only in non-SSL deployments to encrypt credentials passwords.	Disabled 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.
Dispose Resource Transfer	Indicates whether or not system resources allocated for the resource transfer activities should be released. Disabling is not recommended and can only be used for debugging purposes.	Enabled by default.
Log performance data	True indicates that performance data will be logged.	Disabled by default.
Message Bus Notification Enabled	Indicates whether the repository server should interface with the message bus.	Enabled by default.
Modeler Parameter Password Indicator	IBM SPSS 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.
Object Caching: Enable	Indicates whether or not caching of repository objects is enabled. If selected, you need to either enable the internal cache or specify the settings for an external cache.	Disabled by default.

Table 11. Repository configuration options (continued).

Name	Description	Settings
Object Caching: External Cache Address	Address for the external cache. Only needed if the internal cached is not used.	Valid network address.
Object Caching: External Cache Port	Port number for the external cache. Only needed if the internal cached is not used.	Valid port number.
Object Caching: Use Internal	Indicates whether or not the internal cache is used.	Disabled by default.
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.
Remove Deleted Resources	Indicates whether items that are deleted should be removed from the repository. This option should always be selected. It should only be disabled in special cases (for debugging purposes, for example).	Enabled by default.
Repository Maintenance Frequency	Defines frequency (in minutes) for the repository maintenance service. The repository service must be restarted for the changes to take effect.	Integer value. Default is 60 minutes.
Repository Maintenance Master	Defines whether the repository maintenance service should run only on the master node in the server cluster.	True or false. Default is False.
Repository Maintenance Start Date	Defines date and time for the repository maintenance service to start. Invalid dates or dates before the current date are ignored, causing the service to start immediately. If the specified start time is in the past, the service will start at that time the following day.	Date and time in the format *[YYYY-MM-DD] HH:MM:SS*.
Repository Maintenance Start Max	Defines the maximum delay time for the maintenance service to start.	Integer value. Default is 30 minutes.
Repository Maintenance Start Min	Defines the minimum delay time for the maintenance service to start.	Integer value. Default is 5 minutes.
Repository Maintenance Transaction Delay	Defines the percentage for the delay time of the overall maintenance unity or work. For example, if the maintenance transaction delay is 75% (default), and the transaction took 1 second, then it will be followed by a 3 second delay.	Integer value between 1 and 99. Default is 75.

Table 11. Repository configuration options (continued).

Name	Description	Settings
Repository Maintenance Transaction Duration	Defines the duration of each maintenance transaction (in milliseconds) and allows the maintenance services to function without overtaxing system resources and application processing time.	Integer value. Default is 500 milliseconds. A negative value is interpreted as unlimited.
Repository Notification Enabled	Indicates whether the repository server should interface with the notification server.	Disabled by default.
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.	Enabled 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.
Stream Properties Update	If available, this option specifies whether stream properties are updated when the file is published to the repository. Disabling this option can improve performance, and is recommended.	Enabled by default.
Validate Server Executables	Specifies whether or not server executable files should be validated when stored in the repository.	Enabled by default.

Scoring Service

The Scoring Service configuration options allow you to specify settings for scoring.

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 12. Scoring Service configuration options.

Name	Description	Settings
Application Server Authentication for WS-Security	Defines whether to use application server JAAS authentication for WS-Security.	Disabled by default.
Audit Timer Period	The number of milliseconds between audit updates.	Integer value. Default is 3600000.
Default Logging Destination	Default logging destination.	A deployment-specific 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.	Enabled 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 IBM SPSS Collaboration and Deployment Services 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 13. Search configuration options.

Name	Description	Settings
Audit Searches	Log each search in the audit view. See the topic Chapter 15, "Auditing the repository," on page 83 for more information. Note that enabling this option can slow down searches.	Disabled 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.
Search Maintenance Enabled	Defines whether maintenance activities are enabled for the Search Service.	Enabled by default.

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 14. Security configuration options.

Name	Description	Settings
Account Lockout Duration	Number of minutes before automatically unlocking a user who was locked out after exceeding the allowed number of invalid login attempts.	Integer value. Default is 30. A value of 0 means to never unlock users automatically.
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.	Enabled by default.
Cache Session Timeout	Number of minutes before an idle user's security session is removed.	Integer value. Default is 30.
Cached Login Revalidation Interval	Interval in number of minutes to revalidate cached logins. You must restart the server for this setting to take effect.	Integer value. Default is 5.
Disable Clients	Disables login for IBM SPSS Collaboration and Deployment Services client applications (IBM SPSS Collaboration and Deployment Services Deployment Manager, IBM SPSS Collaboration and Deployment Services Deployment Portal, etc.)	Disabled by default.
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.	Enabled by default.
Invalid Login Attempt Count Threshold	Number of failed login attempts to allow before automatically locking out a user.	Integer value. Default is 3. A value of 0 means to never lock out users automatically.
Lowercase User IDs	Forces the internal identifier for a user to be lowercase. This option should be disabled only if a remote user directory depends on case-sensitive user IDs.	Enabled by default.
Message	Message appearing on the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager welcome screen.	Message text. HTML tags can be used to apply formatting.

Table 14. Security configuration options (continued).

Name	Description	Settings
Normalize Principal	Specifies that user names are stored in the database in normalized character format when users are created or imported (*Normalization Form C* as defined by the Unicode technical standard)	Disabled by default.
Resolve Hostnames	Determines whether web service calls should attempt to resolve host names.	Enabled 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 IBM SPSS 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 15. Setup configuration options.

Name	Description	Settings
Log JMS Connection Factory	JNDI name of the log JMS connection factory.	A deployment-specific 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 or server-specific case-sensitive string used by the JNDI service to identify the log JMS queue.
Message Bus JMS Connection Factory	JNDI name of the message bus JMS connection factory.	A deployment-specific 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 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 *https* instead of *http*. Furthermore, the port can be omitted if the server uses the standard *http* port of 80, or the standard *https* port of 443. The server must be restarted for any changes to the prefix to take effect.	URL. **Restriction:** Do not end the URL specification with a slash. For example, specify a value of `http://myserver:8080/myroot` instead of `http://myserver:8080/myroot/`.

CMOR

The CMOR configuration option offers the *UDF Character Limit* setting, allowing you to specify the maximum number of characters that can be passed to database user-defined functions.

The default value is sufficient for most systems and should rarely need to be modified. As a result, the CMOR option is hidden from the standard configuration interface and should only be accessed should errors necessitate increasing the character limit. For example, if the number of characters used in version labels exceeds the specified limit, the system will be unable to retrieve the availableData Provider Definition - Real Time list when selecting a data provider for a scoring configuration and the server log will include truncation errors. If the number of labels cannot be reduced, the UDF character limit needs to be increased. To modify the limit:

1. In the Configuration page, click the **Configuration** link to reveal the hidden settings.
2. In the settings list, under CMOR, click **UDF Character Limit**. The current character limit appears.
3. Modify the value as needed.
4. Click **Set** to establish the new value.
5. Logoff and restart the repository server.

   For some databases, such as SQL Server, DB2, or DB2 on IBM i, the functions cannot be updated automatically to reflect the new value. In these cases, the functions need to be manually updated after shutting down the server but before restarting it as follows:
6. After modifying the configuration value, stop the server.
7. When the server stops, use the existing administration tools for your database to modify the two functions *spsscmor_fn_gl2* and *spsscmor_fn_gl3*. Replace the current character limit value (originally 4000) with the limit specified in the *UDF Character Limit* configuration setting.
8. After updating the values, restart the server.

The following table shows the replacement specifications for each database when increasing the character limit from 4000 to 6000.

Table 16. Example character limit increases.

Database	Old specification	New specification
SQL Server	`@validLabels nvarchar(4000)`	`@validLabels nvarchar(6000)`
DB2	`valid_labels varchar(4000)`	`valid_labels varchar(6000)`
DB2 on IBM i	`valid_labels VARGRAPHIC(4000)`	`valid_labels VARGRAPHIC(6000)`

Chapter 10. MIME types

Multipurpose Internet Mail Extensions, or *MIME*, is a standard for identifying different types of information. MIME originated as an extension of email, 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 IBM SPSS Collaboration and Deployment Services Deployment Manager, accessed by clicking **MIME Types** in the navigation list.

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 IBM SPSS Collaboration and Deployment Services Deployment Portal by default. For external file types (for example, *application/msword*), administrators can assign an icon to the MIME type. See the topic "Adding MIME type mappings" for more information.

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, 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 .

To add a new MIME type mapping:

1. On the MIME Types and File Type Icons page, click **Add New MIME Type**. The Add MIME Types and File Type Icons page appears.
2. 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* .
3. Enter the MIME type being added.
4. Enter the file extensions to associate with the MIME type. Use a space between entries when specifying multiple extensions.

5. 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 needed, select **No**.
6. 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:

1. 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.
2. Modify the settings as necessary. Icons will be changed only if you select a new file or select **No**. To delete an icon, select **No**.
3. 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.

Chapter 11. Reindexing the repository

Indexing is used to optimize IBM SPSS Collaboration and Deployment Services 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. See the topic "Process Management" on page 51 for more information. The repository search is automatically disabled while reindexing is run at startup.

Reindexing can also be performed on demand in the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager by an authorized user. See the topic "Actions" on page 25 for more information.

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 IBM SPSS 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 the repository:

1. In the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager, click **Repository Index** in the navigation list. The Content Repository Indexing page appears.
2. Do one of the following:
   - If no users are logged in to the repository, select **Clear the entire index before reindexing**.
   - If users are still logged in to the 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.

Chapter 12. Repository maintenance

IBM SPSS Collaboration and Deployment Services Repository maintenance can include tasks such as backing up existing data and application settings and cleaning up unused and obsolete data to ensure data integrity and optimal performance.

Over time, the size of the IBM SPSS Collaboration and Deployment Services Repository will tend to get larger. A new object version is stored every time an object is saved. In addition, artifacts created from each job execution accumulate. As a result of this influx of objects and versions, the repository database may grow to a size that can start negatively impacting performance. The performance degradation may result in additional time needed to save a file. In extreme situations, some operations may start much longer than they had in the past or possibly fail with a timeout error. To prevent such problems, periodic removal of unnecessary objects and versions should be performed.

Items that are candidates for removal include the following:

- Unlabeled versions of objects that are not required
- Unneeded Enterprise View versions
- Unnecessary job artifacts
- Expired submitted work. See the topic "Removing expired submitted work" on page 66 for more information.
- Old job histories. See the topic "Managing the job history size" on page 67 for more information.

Deleting unneeded items can be accomplished in a variety of ways. You can identify and remove each item manually. Alternatively, you can use the cleanup utility to perform batch deletion of items that meet specified criteria. Finally, you can use IBM SPSS Collaboration and Deployment Services - Essentials for Python to create automated deletion tasks that can be scheduled for execution at regular intervals. To prevent the deletion of a large number of items from impacting the overall performance of the system, a maintenance service manages the actual deletion.

Repository backup

IBM SPSS Collaboration and Deployment Services Repository data and application setting are stored in a relational database and backup of the repository must performed at the database level with database vendor backup utilities.

Daily database backup is recommended. If necessary, the repository can be reinstalled over a backup copy of the database.

Automatic maintenance service

When you delete an item, the item immediately becomes unavailable to all IBM SPSS Collaboration and Deployment Services Repository clients. However, the item is not removed at that point but is instead flagged for deletion. A maintenance service performs the actual deletion. This service periodically activates and removes flagged items from the system. If all flagged items cannot be removed in the current maintenance window, the items persist in the system until the next service activation. The maintenance service minimizes the impact of deletion tasks on the overall system processing.

There are some exceptions in which items are removed immediately instead of being flagged. If you delete a set of object versions that includes the *LATEST* version, the entire set is deleted immediately to

allow proper reassignment of the *LATEST* label to a new version. Furthermore, performing an export forces all flagged versions to be deleted immediately to prevent deleted items from being included in the export set.

Configuring automatic repository maintenance

The maintenance service performs a variety of tasks, including the following:

- Deleting flagged objects and versions
- Deleting obsolete search indexes
- Removing obsolete job histories
- Removing expired submitted artifacts
- Removing expired pending server connections
- Removing temporary files created during export, import, and promotion activities

The service runs on a schedule defined by a set of configuration parameters. Specify values for these parameters by using the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager. All of the parameters are available in the Repository group of the Configuration page.

1. Select **Repository Maintenance Start Date**. Enter a value indicating the date and time at which the maintenance service should start. Click **Set**.
2. Select **Repository Maintenance Start Max**. Enter a value indicating the longest time period after the specified start time at which the maintenance service should start. If the service is unable to start at the specified time, this is the longest amount of time that the service will attempt to start. Click **Set**.
3. Select **Repository Maintenance Start Min**. Enter a value indicating the shortest time period after the specified start time at which the maintenance service should start. If the service is unable to start at the specified time, this is the shortest amount of time that the service will attempt to start. Click **Set**.
4. Select **Repository Maintenance Frequency**. Enter a value indicating the frequency at which the maintenance service runs. For example, a value of 90 results in the service running every 90 minutes. Click **Set**.
5. Select **Repository Maintenance Transaction Delay**. The overall time for a maintenance transaction consists of the actual maintenance work plus a delay before the next transaction is processed. The delay allows the system to attend to other tasks while the maintenance service is running. Enter a value indicating the percentage of the overall time for a maintenance transaction allocated to this delay. For example, a value of 50% indicates that the transaction work should be followed by a delay equal to the time required to perform the work. In other words, the delay uses half of the total time for the maintenance transaction. Click **Set**.
6. Select **Repository Maintenance Transaction Duration**. Enter a value indicating the time allocated for a maintenance transaction. Click **Set**.
7. If your IBM SPSS Collaboration and Deployment Services server is running in a cluster environment, you can run the maintenance service across all of the cluster nodes or on the master node only. Choose **Repository Maintenance Master** from the Configuration list. Limit the service to the master node by selecting this option. Click **Set**.
8. Restart the IBM SPSS Collaboration and Deployment Services server to begin using the new settings.

For more information about these configuration settings, see “Repository” on page 53.

Removing expired submitted work

Artifacts created in the Submitted Jobs folder automatically expire after a specified number of days, making them visible only to the owner and to administrators. If expired artifacts are not needed beyond their expiration dates, you can configure your system to automatically flag the artifacts for deletion when they expire. When the maintenance service activates, the items will be removed from the repository.

You can control this functionality by using the Configuration page available in the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager.

1. Select **Remove Expired Submitted Artifacts** from the Process Management group.
2. Select the check box to enable this functionality.
3. Click **Set**.

For more information about this configuration setting, see "Process Management" on page 51.

Managing the job history size

Every time a job runs, an entry is added to the job history that details information about that job execution, such as when the execution occurred and what the overall status of the execution was. These entries include references to the job output and to the execution log. If a job runs on a schedule, every execution initiated by the schedule yields a corresponding entry in the job history.

Given that every job execution generates a job history entry, the amount of information being maintained in the job history can become quite large over time. However, some of these history entries may be unneeded. History entries for older executions of a job often become obsolete as newer executions of the job become available. To control the size of the job history, you can define a limit on the number of job history entries to retain for a job version. When the history for a job version exceeds this limit, the oldest history entry becomes obsolete and is removed when the maintenance service activates. For example, if the job history size limit is fifteen, the sixteenth execution results in the first history entry being removed.

You can control this functionality by using the Configuration page available in the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager. To manage the job history entries automatically, perform the following steps:

1. Select **Job History Limit** from the Process Management group. Enter the number of history entries to retain for each job version. Click **Set**.
2. Select **Remove Obsolete Job Histories** in the Process Management group. Select the check box to enable removal of the oldest job version histories in excess of the job history limit. Click **Set**.

For more information about these configuration settings, see "Process Management" on page 51.

Monitoring maintenance activities

Maintenance service activity summaries can be included in the system log files, enabling you to identify the tasks performed when the service activates.

To enable maintenance service logging:

1. Open the *log4j.xml* configuration file in a text editor.
2. Locate the category element for the *com.spss.process.internal.maintenance* logger.
3. Set the logging level for this logger to *DEBUG*.
4. Save your changes.

When the maintenance service activates, the following messages will be added to the log output:

- Removed *N* expired submitted executions in the time allotted.
- Removed *N* obsolete executions in the time allotted.

For more information about logging services, see the *Installation and Configuration Guide*.

Batch deletion

Deleting a large number of items can be extremely tedious if you need to add each item separately. However, if the items share a set of characteristics, you can use the cleanup utility to identify, select, and delete items in bulk.

To use this utility, you specify the criteria that must be matched for an item to be selected and deleted. The selection criteria can be based on the following characteristics:

- folder
- MIME type
- Label presence
- Number of versions
- Creation date

For example, you can use the cleanup utility to delete all but the last three versions of every IBM SPSS Statistics syntax file in a specified folder. Alternatively, you can delete all unlabeled version of IBM SPSS Modeler streams older than a specified date.

If the automatic maintenance framework is enabled, the selected items are flagged for subsequent deletion at the next available opportunity. If the maintenance framework is disabled, the items are immediately deleted.

The cleanup utility is entirely Java-based and runs on any supported IBM SPSS Collaboration and Deployment Services platform. The utility is available in the following folder:

```
<repository install path>/applications/cleanup
```

Note that item deletion is permanent; once an item is deleted, it cannot be recovered. To avoid unnecessary risk, consider backing up the data before deleting files with this utility.

You can execute the cleanup utility from the command line or create job steps for automatic, recurring processing.

It is recommended to back up the repository database before deleting files with this utility. Alternatively, you can use the IBM SPSS Collaboration and Deployment Services export facility to create a backup of any folders that will be processed by the cleanup utility.

Running the cleanup utility

The command for running the cleanup utility has the following structure:

```
cleanup <parameter=value parameter=value ...>
```

The `cleanup` command is followed by a space-delimited list of parameters and their values that define the deletion task. Each parameter specification includes the parameter name, an equals sign, and the parameter value. The Table 17 table describes each parameter.

Table 17. Cleanup utility parameters.

Parameter	Use	Description
`connectionURL`	Required	The IBM SPSS Collaboration and Deployment Services RepositoryURL
`userid`	Required	A valid native IBM SPSS Collaboration and Deployment Services user identifier for connecting to the repository server. The user must have sufficient permissions for deleting any selected items. Typically the identifier corresponds to an administrator.

Table 17. Cleanup utility parameters (continued).

Parameter	Use	Description
`password`	Required	The password for the specified user
`resource`	Required	The path to a repository folder or file. This parameter may be specified multiple times.
`includeSubFolders`	Optional	A boolean value indicating whether or not subfolders should be searched. Default is false.
`includeType`	Optional	MIME types of objects to include. The comparison is not case-sensitive, but must match the exact text. This value may be specified multiple times. Default is all types.
`excludeType`	Optional	MIME types of objects to exclude. The comparison is not case-sensitive, but must match the exact text. This value may be specified multiple times. Default is no exclusions.
`deleteLabeled`	Optional	A boolean value indicating whether or not labeled versions should be deleted. Default is false.
`versionsToKeep`	Optional	The number of most recent versions that should be preserved. Default is 0.
`olderThan`	Optional	Only resources created before the specified date are selected. Times are localized to the machine running the cleanup utility for comparison. Default is no date filter.
`logfile`	Optional	The path to a local file that will be used for logging the results. Default is no log file.
`testMode`	Optional	A boolean value indicating whether or not the selected items should be deleted. A value of *true* results in the objects/versions being selected without actually being deleting. Default is false.

You invoke the cleanup utility by using the following steps:

1. Verify that the system *Path* environment variable includes your Java path.
2. At a command prompt, navigate to the directory containing the cleanup utility.
3. Type `cleanup`, followed by the list of parameters and values that define your deletion task.
4. Entering the command initiates the task.

For example, the following command recurses all subfolders in the */CleanupData* folder, selecting unlabeled versions for deletion. The `testMode` parameter prevents the versions from actually being deleted, allowing you to review the *cleanup.log* file to identify the selected versions that would be deleted if you removed `testMode`.

```
cleanup userid=admin password=pass connectionURL=http://localhost:8080
    testMode=true resource=/CleanupData includeSubFolders=true logfile=cleanup.log
```

Creating batch deletion jobs

You can initiate batch deletion from a IBM SPSS Collaboration and Deployment Services job using a General job step.

To create a job step for batch deletion in IBM SPSS Collaboration and Deployment Services Deployment Manager, perform the following steps:

1. Add a General job step to a job.
2. Click the job step to modify the properties.
3. On the General tab, type a name for the step. For the **Command To Run**, type the full path to the cleanup utility for your system followed by the cleanup utility parameters defining the deletion task.

4. If the deletion task includes the `logfile` parameter and you want the log to be saved to the IBM SPSS Collaboration and Deployment Services Repository, use the Output Files tab to specify the target location for the file.
5. Save the job.

The job can be executed manually as needed, or you can create a schedule that automatically runs the job at specified times or in response to system events. For more information on General job steps and scheduling jobs, see the IBM SPSS Collaboration and Deployment Services Deployment Manager documentation.

Chapter 13. Notifications

IBM SPSS Collaboration and Deployment Services provides the mechanisms of *notifications* and *subscriptions* for keeping the users informed about changes to IBM SPSS Collaboration and Deployment Services Repository objects and job processing results. Both notifications and subscriptions generate email messages when corresponding events occur. For example, when a job fails, IBM SPSS Collaboration and Deployment Services can automatically send an email 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 email that is sent to any recipients associated with the event.

Notification templates included in the default repository installation can be found 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.xsl* and *job_failure.xsl*, correspond to the success and failure of job executions. If a job completes successfully, IBM SPSS 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 template structure

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 email processor. The following sample template generates a simple, generic email message indicating the success of the corresponding job.

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

For more information about Velocity templates, see the Apache Velocity project documentation.

Message properties

Email 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 email address assigned to the message. Default SMTP properties are listed under repository notification configuration options. If the Sun JVM is used with the repository installation, SMTP properties will correspond to the JavaMail API properties for message handling defined in the following table. Note that these properties may be different for different Java environments. For detailed information about SMTP properties, see the JVM vendor documentation.

Table 18. 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.

Table 18. Message properties (continued).

Message Property	Attribute	Event Property	Description
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 email 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.
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: • `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.

```
<html>
<head>
<meta http-equiv='Content-Type' content='text/html;charset=utf-8'/>
</head>
<body>
<p>The job <b>${JobName}</b> started ${JobStart} and #if($JobSuccess) completed successfully #else failed #end ${JobEnd}.

<p>To review the job log, go to <a href='${JobStatusURL}'>${JobStatusURL}</a>.</p>
```

```
<hr><p>This is a machine-generated message. Please do not reply directly. If you do not wish to receive this notification,
remove yourself from the notification list or contact your Repository administrator.</p>
</body>
</html>
```

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. IBM SPSS Collaboration and Deployment Services jobs cannot be opened outside IBM SPSS Collaboration and Deployment Services 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 a IBM SPSS 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 IBM SPSS Deployment
Services administrator.</a></p></body>
</html>
```

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 `$` notation to insert event property variables into the message subject. See the topic "Message content" on page 73 for more information.
3. Define the MIME type of the message. The MIME type value is specified in the square brackets following `messageContent`. For a plain text message, use a value of *text/plain*. For an HTML message, use a value of *text/html*. See the topic "Message format" on page 74 for more information.
4. Modify the value assigned to `messageContent`. Use the `$` 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.

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.

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.

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.**
- To expand the details for a job step, which include job step log and any resulting output, click **+** next to the job step name.

  The following information is presented for a job step:

  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 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.**

Job status

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

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.

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.

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.**
- To expand the details for a job step, which include job step log and any resulting output, click **+** next to the job step name.

  The following information is presented for a job step:

  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 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.**

Optimizing notification service performance

The overall performance of the notification service is a combination of the performance of IBM SPSS 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 IBM SPSS 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 (email 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 email 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 load 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 IBM SPSS 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 page 47

Dedicated SMTP server

The performance of the delivery channel, such as an email server, is the critical factor controlling the overall performance of the notification service. For IBM SPSS Collaboration and Deployment Services notifications, it is strongly recommended to use a fast, dedicated SMTP server rather than the regular corporate email 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 email server on the same host as the repository, which reduces the time it takes notification service to communicate with the email 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 IBM SPSS 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 IBM SPSS 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 use available resources and handle delivery of the notification messages much more efficiently.

Minimize the number of custom templates.

IBM SPSS 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. See the topic "Editing notification templates" on page 75 for more information. 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 often sufficient to specify a single subscription for the parent folder and use the `Apply to Subfolders` option. For more information, see the IBM SPSS Collaboration and Deployment Services 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 IBM SPSS 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 IBM SPSS Collaboration and Deployment Services 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 email server and the default sender's email address during installation of the repository, additional email configuration is not usually required in order for IBM SPSS 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. See the topic "Notification" on page 47 for more information.

The IBM SPSS 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: IBM SPSS Deployment Services: New version of ChurnAnalysis created
      Sent: 4/5/2010 2:35 PM

The following recipient(s) could not be reached:

      jsmiht@mycompany.com on 4/5/2010 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 email 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. Notification and subscription delivery failures are logged in repository auditing views. See the topic Chapter 15, "Auditing the repository," on page 83 for more information.
- To determine the cause of the notification failure, it is recommenced to enable the debugging mode. See the topic "Debugging the notification service" on page 79 for more information.
- **`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.

Chapter 14. JMS configuration for process management

IBM SPSS Collaboration and Deployment Services uses Java Messaging Service (JMS) to communicate with third-party applications and trigger job processing based on IBM SPSS Collaboration and Deployment Services 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 and 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 provider.

The application server JMS settings can be modified to increase concurrency limits when IBM SPSS Collaboration and Deployment Services performance must be optimized, for example, when a high number of jobs are processed concurrently. For information on increasing JMS concurrency limit, see the topic below. This chapter also provides an example demonstrating how to set up job processing based on repository events.

Increasing JMS concurrency limits

When IBM SPSS Collaboration and Deployment Services performance must be optimized due to high workload, for example, a lot of jobs running concurrently, it may be necessary to modify the application server JMS setting to increase concurrency limits. The following are general steps for WebSphere, JBoss, and WebLogic. For more detailed information, consult the application server documentation.

WebSphere

1. In WebSphere Integrated Solutions Console select
   Resources > **JMS** > **Activation Specifications**
2. Open **CaDSProcessEventActivationSpec** and increase the value of `Maximum concurrent MDB invocations per endpoint`.
3. Restart the server.

JBoss

1. Increase the value of **MaximumSize** element in `<JBoss server directory>/conf/standardjboss.xml`.
   In the following example the value of **MaximumSize** is set to 150 (default is 15).

```
<invoker-proxy-binding>
      <name>message-driven-bean</name>
      <invoker-mbean>default</invoker-mbean>
      <proxy-factory>org.jboss.ejb.plugins.jms.JMSContainerInvoker</proxy-factory>
      <proxy-factory-config>
        <JMSProviderAdapterJNDI>DefaultJMSProvider</JMSProviderAdapterJNDI>
        <ServerSessionPoolFactoryJNDI>StdJMSPool</ServerSessionPoolFactoryJNDI>
        <CreateJBossMQDestination>true</CreateJBossMQDestination>
        <!-- WARN: Don't set this to zero until a bug in the pooled executor is fixed -->
        <MinimumSize>1</MinimumSize>
        <MaximumSize>150</MaximumSize>
        <KeepAliveMillis>30000</KeepAliveMillis>
        <MaxMessages>1</MaxMessages>
        <MDBConfig>
          <ReconnectIntervalSec>10</ReconnectIntervalSec>
          <DLQConfig>
            <DestinationQueue>queue/DLQ</DestinationQueue>
            <MaxTimesRedelivered>200</MaxTimesRedelivered>
            <TimeToLive>0</TimeToLive>
```

```
        </DLQConfig>
      </MDBConfig>
    </proxy-factory-config>
  </invoker-proxy-binding>
```

2. Restart the server. The change will affect all deployed message-driven beans.

WebLogic

Use a WebLogic work manager to control the number of active threads.

1. Create a new work manager and target it to the WebLogic server used to run IBM SPSS Collaboration and Deployment Services.
2. Update the deployment descriptor to reference the new work manager.
3. Modify `weblogic-ejb-jar.xml` in `process-ejb.jar`, which can be found in `<repository installation directory>/platform/deployables/process-ejb.ear`. Append the following:

```
<dispatch-policy>PASWWorkManager</dispatch-policy>
<weblogic-enterprise-bean>
  <ejb-name>ProcessEventMDB</ejb-name>
   <message-driven-descriptor>
    <pool>
     <max-beans-in-free-pool>20</max-beans-in-free-pool>
     <initial-beans-in-free-pool>1</initial-beans-in-free-pool>
    </pool>
    <destination-jndi-name>queue/SPSSProcess</destination-jndi-name>
    <connection-factory-jndi-name>ProcessConnectionFactory</connection-factory-jndi-name>
   </message-driven-descriptor>
 <dispatch-policy>PASWWorkManager</dispatch-policy>
</weblogic-enterprise-bean>
```

4. Update `process-ejb.ear` on the application server and adjust its settings in the administration console.

Message-based processing example

Message-based scheduling functionality of IBM SPSS 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 IBM SPSS Modeler stream used in one of the job steps is updated.

The procedure involves the following steps:

1. Using IBM SPSS Collaboration and Deployment Services Deployment Manager, create a JMS message domain.
2. 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 IBM SPSS Modeler stream as in the following example:

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

   The repository resource ID of the IBM SPSS Modeler stream can be found in the object properties.
3. Set up a notification for the IBM SPSS Modeler stream based on the JMS subscriber you have defined.
4. To test the message-based schedule, the stream must be opened in IBM SPSS Modeler, modified, and stored in the repository. If everything has been set up correctly, the schedule will trigger the job.

Chapter 15. 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, the IBM SPSS Collaboration and Deployment Services 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 IBM SPSS Collaboration and Deployment Services 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

The repository provides several database tables for recording system events and changes to objects. When the 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, BIRT Report Designer for IBM SPSS, included in the IBM SPSS 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.

The 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 IBM SPSS 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 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 the 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 IBM SPSS Modeler streams, IBM SPSS 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 IBM SPSS Modeler stream, an IBM SPSS 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.

NEXTSCHEDULEDTIME. 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 IBM SPSS 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.

ATTRIBUTECATEGORICALVALUE. An allowable value for the attribute for multivalue attributes.

NUMERICALUPPERBOUND. The upper bounds value allowable for numerical attributes.

NUMERICALLOWERBOUND. The lower bounds value allowable for numerical attributes.

Stream node (SPSSPLAT_V_STREAMNODE)

The Stream Node view presents the information for the nodes in IBM SPSS 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 FileOuput 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

IBM SPSS Collaboration and Deployment Services also provides database facilities for logging the operations of the services for IBM SPSS Collaboration and Deployment Services - Scoring. 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 the SPSSSCORE_LOG table. Each row in the table corresponds to a scoring service request.

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. See the topic "XML schema" on page 95 for more information. The same information is available in relational format from the scoring log view.

Clean-up 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 before 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 IBM SPSS Collaboration and Deployment Services 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 IBM SPSS 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 IBM SPSS Collaboration and Deployment Services 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 Data Provider Definition - Real Time. 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 DPD output (SPSSSCORE_V_LOG_DPD_OUTPUT)

This view contains the information about the data that was passed to the scoring service from a Data Provider Definition - Real Time. If a Data Provider Definition - Real Time is not being used, these entries are absent.

SERIAL. The unique identifier of the scoring request row.

DO_TABLE. The name of the table used in the Data Provider Definition - Real Time.

DO_ROW. The row number of the Data Provider Definition - Real Time data row starting at 1.

DO_NAME. The name of an input field corresponding to the name of the column in the Data Provider Definition - Real Time.

DO_VALUE. Input value.

Scoring request input (SPSSSCORE_V_LOG_REQUEST_INPUT)

This view contains the information about the data used as input for the scoring service request.

SERIAL. The unique identifier of the scoring request row.

RI_TABLE. The name of the table used in the request.

RI_ROW. The row number of the request data row starting at 1.

RI_NAME. The name of an input field corresponding to the name of the column in the request.

RI_VALUE. Input value.

Scoring request properties (SPSSSCORE_V_LOG_REQUEST_PROP)

This view contains the information about the properties associated with an input table.

SERIAL. The unique identifier of the scoring request row.

RI_TABLE. The name of the table used in the request.

RI_PROP_NAME. The name of the property.

RI_PROP_VALUE. The value for the property.

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 in the SPSSSCORE_V_LOG_PROPERTY view for each row in the 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.

XML schema

The following XML schema is registered with the database and used for the `INFO` column of the `SPSSSCORE_LOG` table. This schema is required for MS SQL Server and Oracle. It is not required on DB2.

```
<?xml version="1.0" encoding="UTF-8"?>
<xs:schema
    attributeFormDefault="unqualified"
    elementFormDefault="qualified"
    targetNamespace="http://xml.spss.com/scoring/logging"
    version="2.0"
    jaxb:version="2.0"
    xmlns:jaxb="http://java.sun.com/xml/ns/jaxb"
    xmlns:spss_ss_logging="http://xml.spss.com/scoring/logging"
    xmlns:xs="http://www.w3.org/2001/XMLSchema">

    <!-- ************************************************************** -->
    <!-- SIMPLE TYPES                                                   -->
    <!-- ************************************************************** -->
    <xs:simpleType name="pevDataType">
        <xs:annotation>
            <xs:documentation>The type of this column. This maps to the same types defined by
              the DD EventServer. We will map these types to the SQL types using the same
              mapping that the DD Event Server uses.</xs:documentation>
        </xs:annotation>
        <xs:restriction base="xs:NMTOKEN">
            <xs:enumeration value="boolean"/>
            <!--    <xs:enumeration value="character"></xs:enumeration> not needed, as string
                    should be sufficient for mapping to SQL -->
            <xs:enumeration value="date"/>
            <xs:enumeration value="daytime"/>
            <xs:enumeration value="decimal"/>
            <xs:enumeration value="double"/>
            <xs:enumeration value="float"/>
            <xs:enumeration value="integer"/>
            <xs:enumeration value="long"/>
            <xs:enumeration value="string"/>
            <xs:enumeration value="timestamp"/>
        </xs:restriction>
    </xs:simpleType>

    <xs:attributeGroup name="nillableValueAttributeGroup">
        <xs:attribute name="value" type="xs:string" use="optional">
            <xs:annotation>
                <xs:documentation>A value, in string representation.  If this attribute is not
                specified, the value is considered to be null.  The text representation of the
                numeric types is obvious, but several types are not.  The format of the
                non-numeric types must be as follows: boolean='true'(case insensitive) or '1'
                or 'false'(case insensitive) or '0', date='yyyy-MM-dd', daytime='HH:mm:ss', and
                timestamp='yyyy-MM-ddTHH:mm:ss'.</xs:documentation>
            </xs:annotation>
        </xs:attribute>
    </xs:attributeGroup>

    <!-- ************************************************************** -->
    <!-- COMPLEX TYPES                                                  -->
```

```
<!-- ************************************************************* -->
<xs:complexType name="modelInputValue">
    <xs:annotation>
        <xs:documentation>This element is optionally returned as part of the scoreResult
        element.  If the configuration is programmed to return the model input fields
        (see spss_ss:modelInputMetadataField), then this element contains the value that
        was used to produce the score.  The value might be null.</xs:documentation>
    </xs:annotation>
    <xs:attribute name="name" type="xs:string" use="required">
        <xs:annotation>
            <xs:documentation>The name of the input item.</xs:documentation>
        </xs:annotation>
    </xs:attribute>
    <xs:attribute name="type" type="spss_ss_logging:pevDataType" use="required">
        <xs:annotation>
            <xs:documentation>The data type of the input item.</xs:documentation>
        </xs:annotation>
    </xs:attribute>
    <xs:attributeGroup ref="spss_ss_logging:nillableValueAttributeGroup"/>
</xs:complexType>

<xs:complexType name="inputTable">
    <xs:annotation>
        <xs:documentation>One table of input values, may contain zero or more
        rows.</xs:documentation>
    </xs:annotation>
    <xs:sequence>
        <xs:element name="columns" type="spss_ss_logging:inputColumn" minOccurs="1"
        maxOccurs="unbounded">
            <xs:annotation>
                <xs:documentation>An ordered list of column names</xs:documentation>
            </xs:annotation>
        </xs:element>
        <xs:element name="rowValues" type="spss_ss_logging:rowValues" minOccurs="0"
        maxOccurs="unbounded">
            <xs:annotation>
                <xs:documentation>A row of values, value order must match defined column
                order.</xs:documentation>
            </xs:annotation>
        </xs:element>
    </xs:sequence>
    <xs:attribute name="sourceTable" type="xs:string" use="required">
        <xs:annotation>
            <xs:documentation>This attribute holds the name of the source table as defined
            in the model.</xs:documentation>
        </xs:annotation>
    </xs:attribute>
</xs:complexType>

<xs:complexType name="inputColumn">
    <xs:annotation>
        <xs:documentation>Describes a column in the designated input table.  If the
        configuration is programmed to return the model input fields (see
        spss_ss:modelInputMetadataField), then this element contains the value that
        was used to produce the score.  The value might be null.</xs:documentation>
    </xs:annotation>
    <xs:attribute name="name" type="xs:string" use="required">
        <xs:annotation>
            <xs:documentation>The name of the input item.</xs:documentation>
        </xs:annotation>
    </xs:attribute>
    <xs:attribute name="type" type="spss_ss_logging:pevDataType" use="required">
        <xs:annotation>
            <xs:documentation>The data type of the input item.</xs:documentation>
        </xs:annotation>
    </xs:attribute>
</xs:complexType>

<xs:complexType name="inputTableWithProperties" >
    <xs:annotation>
        <xs:documentation>Input tables can have loggable properties</xs:documentation>
    </xs:annotation>
    <xs:complexContent>
        <xs:extension base="spss_ss_logging:inputTable">
            <xs:sequence>
                <xs:element name="RequestInputProperties"
                type="spss_ss_logging:requestInputProperties" minOccurs="0" maxOccurs="1">
                    <xs:annotation>
                        <xs:documentation>Properties that are associated with an input
                        table</xs:documentation>
                    </xs:annotation>
                </xs:element>
            </xs:sequence>
        </xs:extension>
```

```
    </xs:complexContent>
</xs:complexType>

<xs:complexType name="requestInputProperties">
    <xs:annotation>
        <xs:documentation>Properties that are associated with an input table</xs:documentation>
    </xs:annotation>
    <xs:sequence>
        <xs:element name="property" type="spss_ss_logging:nameValueType" minOccurs="1"
        maxOccurs="unbounded">
            <xs:annotation>
                <xs:documentation>Properties that are associated with an input
                table</xs:documentation>
            </xs:annotation>
        </xs:element>
    </xs:sequence>
</xs:complexType>

<xs:complexType name="columnNames">
    <xs:annotation>
        <xs:documentation/>
    </xs:annotation>
    <xs:sequence>
        <xs:element name="name" type="xs:string" minOccurs="1" maxOccurs="unbounded"/>
    </xs:sequence>
</xs:complexType>

<xs:complexType name="rowValues">
    <xs:annotation>
        <xs:documentation>One row of values, note that a value may be nill.</xs:documentation>
    </xs:annotation>
    <xs:sequence>
        <xs:element name="value" type="spss_ss_logging:nillableValue" minOccurs="1"
        maxOccurs="unbounded"/>
    </xs:sequence>
</xs:complexType>

<xs:complexType name="output">
    <xs:sequence>
        <xs:element name="columnNames" type="spss_ss_logging:columnNames">
            <xs:annotation>
                <xs:documentation>An ordered list of column names</xs:documentation>
            </xs:annotation>
        </xs:element>
        <xs:element name="rowValues" type="spss_ss_logging:rowValues" minOccurs="1"
        maxOccurs="unbounded">
            <xs:annotation>
                <xs:documentation>A row of score data, following the order in the
                columnNames element</xs:documentation>
            </xs:annotation>
        </xs:element>
    </xs:sequence>
</xs:complexType>

<xs:complexType name="nameValueType">
    <xs:annotation>
        <xs:documentation>A name value pair.</xs:documentation>
    </xs:annotation>
    <xs:attribute name="name" type="xs:string" use="required"/>
    <xs:attribute name="value" type="xs:string" use="required"/>
</xs:complexType>

<xs:complexType name="context">
    <xs:annotation>
        <xs:documentation>This element contains all the context data inputs to the score
        request.</xs:documentation>
    </xs:annotation>
    <xs:sequence>
        <xs:element name="columnNames" type="spss_ss_logging:columnNames">
            <xs:annotation>
                <xs:documentation>An ordered list of column names</xs:documentation>
            </xs:annotation>
        </xs:element>
        <xs:element name="rowValues" type="spss_ss_logging:rowValues" minOccurs="1"
        maxOccurs="unbounded">
            <xs:annotation>
                <xs:documentation>A row of context data, following the order in the
                columnNames element</xs:documentation>
            </xs:annotation>
        </xs:element>
    </xs:sequence>
    <xs:attribute name="table" type="xs:string" use="required">
        <xs:annotation>
            <xs:documentation>This attribute describes which context table the input data
```

```
            belongs to.</xs:documentation>
        </xs:annotation>
    </xs:attribute>
</xs:complexType>

<xs:complexType name="nillableValue">
    <xs:annotation>
        <xs:documentation>Nillable elements and simpleTypes are not well supported by most
        of the popular frameworks, especially Castor.  Instead of a nillable string element,
        use an optional string attribute to represent null values.</xs:documentation>
    </xs:annotation>
    <xs:attributeGroup ref="spss_ss_logging:nillableValueAttributeGroup"/>
</xs:complexType>

<!-- ************************************************************** -->
<!-- ELEMENTS                                                       -->
<!-- ************************************************************** -->
<xs:element name="Info">
    <xs:complexType>
        <xs:sequence>
            <xs:element name="Output" type="spss_ss_logging:output" minOccurs="0" maxOccurs="1">
                <xs:annotation>
                    <xs:documentation>PA has the ability to generate multiple outputs
                    (multiple offers).  There will be one OutputRow for each output
                    (for each offer). </xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="Input" type="spss_ss_logging:modelInputValue" minOccurs="0"
            maxOccurs="unbounded">
                <xs:annotation>
                    <xs:documentation>This might be a PEV AV input column or data from some
                    other source. THIS ELEMENT IS NOW DEPRECATED. The DpdOutputs and
                    RequestInputs elements will be used for all future log entries, and this
                    element is kept for backwards compatibility.</xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="ContextInput" type="spss_ss_logging:context" minOccurs="0"
            maxOccurs="unbounded">
                <xs:annotation>
                    <xs:documentation>Context data that is fed into the DPD (data engine)
                    and not necessarily into the model.  </xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="DpdOutputs" type="spss_ss_logging:inputTable" minOccurs="0"
            maxOccurs="unbounded">
                <xs:annotation>
                    <xs:documentation>Zero to N input tables. The data contained in each
                    table represents the values that have been provided by the data service.
                    If a RT-DPD is not used for the model, these entries are not
                    present.</xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="RequestInputs" type="spss_ss_logging:inputTableWithProperties"
            minOccurs="0" maxOccurs="unbounded">
                <xs:annotation>
                    <xs:documentation>Zero to N score request input tables. The data
                    contained in each table represents the inputs provided with the score
                    request.</xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="Metric" type="spss_ss_logging:nameValueType" minOccurs="0"
            maxOccurs="unbounded">
                <xs:annotation>
                    <xs:documentation>A metric which is defined by either the HSS engine
                    or the provider.
                    Value is a double represented as a string to account for the
                    correct precision and scale.
                    An example might be the time to produce the output.</xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="Property" type="spss_ss_logging:nameValueType" minOccurs="0"
            maxOccurs="unbounded">
                <xs:annotation>
                    <xs:documentation>A property value.  The name is the name of the
                    property.</xs:documentation>
                </xs:annotation>
            </xs:element>
            <xs:element name="ModelObjectId" type="xs:string" minOccurs="1" maxOccurs="1"/>
            <xs:element name="ModelVersionMarker" type="xs:string" minOccurs="1" maxOccurs="1"/>
            <xs:element name="ConfigurationName" type="xs:string" minOccurs="1" maxOccurs="1"/>
            <xs:element name="ModelInputTable" type="xs:string" minOccurs="0" maxOccurs="1">
                <xs:annotation>
                    <xs:documentation>THIS ELEMENT IS NOW DEPRECATED.</xs:documentation>
                </xs:annotation>
```

```
            </xs:element>
          </xs:sequence>
        </xs:complexType>
    </xs:element>
</xs:schema>
```

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'
```

Chapter 16. nativestore schema reference

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

Example XML

```
<?xml version="1.0" encoding="UTF-8"?>
<nativestore xmlns="spssnative">
  <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

Root element for importing local users and their groups into IBM SPSS Collaboration and Deployment Services.

Child elements

user, obsolete

user element

User to be added or updated.

Parent element

nativestore

Child elements

group, role

Attributes

Table 19. Attributes for the user element.

Name	Type	Use	Default	Description
`userID`	string	required	*no default value*	User ID that will be used to log in to the system.

Table 19. Attributes for the user element (continued).

Name	Type	Use	Default	Description
password	string	optional	*no default value*	Usually a plain-text password. If the encrypted 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 *not* exposed in the IBM SPSS Collaboration and Deployment Services user interfaces.
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 encrypted attribute is usually omitted.

Example XML

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

group element

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

Type: string

Parent element

user

Example XML

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

role element

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

Type: string

Parent element

user

obsolete element

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.

Parent element

nativestore

Child elements

user, group

Example XML

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

user element

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

Type: string

Parent element

obsolete

Example XML

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

group element

Group name to be removed.

Type: string

Parent element

obsolete

Example XML

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

Notices

This information was developed for products and services offered in the U.S.A.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing
IBM Corporation
North Castle Drive
Armonk, NY 10504-1785
U.S.A.

For license inquiries regarding double-byte (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing
Legal and Intellectual Property Law
IBM Japan Ltd.
1623-14, Shimotsuruma, Yamato-shi
Kanagawa 242-8502 Japan

The following paragraph does not apply to the United Kingdom or any other country where such provisions are inconsistent with local law: INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some states do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM Web sites are provided for convenience only and do not in any manner serve as an endorsement of those Web sites. The materials at those Web sites are not part of the materials for this IBM product and use of those Web sites is at your own risk.

IBM may use or distribute any of the information you supply in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Software Group
ATTN: Licensing
200 W. Madison St.
Chicago, IL; 60606
U.S.A.

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

Any performance data contained herein was determined in a controlled environment. Therefore, the results obtained in other operating environments may vary significantly. Some measurements may have been made on development-level systems and there is no guarantee that these measurements will be the same on generally available systems. Furthermore, some measurements may have been estimated through extrapolation. Actual results may vary. Users of this document should verify the applicable data for their specific environment.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

All statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to the names and addresses used by an actual business enterprise is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at "Copyright and trademark information" at www.ibm.com/legal/copytrade.shtml.

Adobe, the Adobe logo, PostScript, and the PostScript logo are either registered trademarks or trademarks of Adobe Systems Incorporated in the United States, and/or other countries.

Intel, Intel logo, Intel Inside, Intel Inside logo, Intel Centrino, Intel Centrino logo, Celeron, Intel Xeon, Intel SpeedStep, Itanium, and Pentium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

Linux is a registered trademark of Linus Torvalds in the United States, other countries, or both.

Microsoft, Windows, Windows NT, and the Windows logo are trademarks of Microsoft Corporation in the United States, other countries, or both.

UNIX is a registered trademark of The Open Group in the United States and other countries.

Java and all Java-based trademarks and logos are trademarks or registered trademarks of Oracle and/or its affiliates.

Other product and service names might be trademarks of IBM or other companies.

Index

F

G

H

I

J

K

L

M

N

O

V

W

X

Printed in USA