

IBM SPSS Collaboration and Deployment Services  
Deployment Manager  
Version 7 Release 0

*User's Guide*

**IBM**

**Note**

Before using this information and the product it supports, read the information in "Notices" on page 261.

**Product Information**

This edition applies to version 7, release 0, modification 0 of IBM SPSS Collaboration and Deployment Services and to all subsequent releases and modifications until otherwise indicated in new editions.

© Copyright IBM Corporation 2000, 2015.

US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

---

# Contents

## Chapter 1. Overview . . . . . 1

IBM SPSS Collaboration and Deployment Services . . . . .	1
Collaboration . . . . .	1
Deployment . . . . .	2
System architecture . . . . .	2
IBM SPSS Collaboration and Deployment Services Repository . . . . .	3
IBM SPSS Collaboration and Deployment Services Deployment Manager . . . . .	4
IBM SPSS Collaboration and Deployment Services Deployment Portal . . . . .	4
IBM SPSS Collaboration and Deployment Services Enterprise View . . . . .	5
Execution servers . . . . .	5
Scoring server . . . . .	6
BIRT Report Designer for IBM SPSS . . . . .	6

## Chapter 2. What is new in this release. . . . . 7

What is new for IBM SPSS Collaboration and Deployment Services Deployment Manager users . . . . .	7
--	---

## Chapter 3. Getting started . . . . . 9

Launching IBM SPSS Collaboration and Deployment Services Deployment Manager . . . . .	9
Navigating through the system . . . . .	9
Using the mouse versus pressing Enter . . . . .	9
Dragging and dropping items in the user interface . . . . .	9
Accessing help . . . . .	9
Entry field content assist . . . . .	10
Field conventions . . . . .	10
Naming conventions within the system . . . . .	10
Exiting the system . . . . .	11

## Chapter 4. The Content Explorer. . . . . 13

Content Explorer overview . . . . .	13
What Is a content object? . . . . .	13
Organization of the Content Explorer . . . . .	13
Working with servers . . . . .	14
Creating new Content Server connections . . . . .	14
Logging in to a server . . . . .	15
Logging off from a server . . . . .	15
Changing server passwords . . . . .	15
Working with files . . . . .	16
Opening external files . . . . .	16
Permissions inheritance . . . . .	16
Working with the repository . . . . .	16
Adding files to the repository . . . . .	16
Downloading files from the repository . . . . .	17
Deleting files from the repository . . . . .	17
Searching . . . . .	17
Accessing the search dialogs . . . . .	18
Simple Search . . . . .	18
Advanced Search . . . . .	18
Viewing Search Results . . . . .	21
Terms Not Included in a Search . . . . .	21

Object Locking . . . . .	21
Locking Objects . . . . .	22
Viewing the Locked Objects Table . . . . .	22
Unlocking Objects . . . . .	22
Obtaining updates . . . . .	23

## Chapter 5. Properties . . . . . 25

Working with object properties . . . . .	25
Viewing object properties . . . . .	25
Editing object properties . . . . .	26
Working With Expiration Dates and Expired Files . . . . .	31
Working with server properties and user preferences . . . . .	33
Server properties . . . . .	33
User preferences . . . . .	35
Version properties . . . . .	37
Working with custom properties . . . . .	38
Verifying access privileges to create custom properties . . . . .	38
Accessing the custom properties dialog box . . . . .	38
Creating custom properties . . . . .	39
Editing custom properties . . . . .	40
Searching for custom properties . . . . .	40
Deleting custom properties . . . . .	41
Working with topics . . . . .	41
Process overview . . . . .	42
Verifying access privileges to create topic definitions . . . . .	42
Working with topic definitions . . . . .	42
Searching for topics . . . . .	44
Bulk property updates . . . . .	45
Updating general properties in bulk . . . . .	45
Updating version properties in bulk . . . . .	45
Updating permissions in bulk . . . . .	46

## Chapter 6. Resource definitions . . . . . 47

Credentials . . . . .	47
Adding new credentials . . . . .	47
Server Process Credential . . . . .	48
Data sources . . . . .	48
Selecting a data source definition type . . . . .	49
Specifying a DSN for an ODBC data source . . . . .	49
Specifying a JDBC name and URL . . . . .	49
Specifying a JNDI name for an application server data source . . . . .	52
Specifying properties for data service data sources . . . . .	52
Modifying data source definitions . . . . .	54
Message domains . . . . .	54
Creating a new message domain . . . . .	54
Modifying message domain definitions . . . . .	55
Promotion policies . . . . .	55
Adding promotion policies . . . . .	55
Modifying promotion policies . . . . .	57
Deleting promotion policies . . . . .	57
Server definitions . . . . .	57

Adding new server definitions . . . . .	58
Modifying server definitions. . . . .	59
Server clusters . . . . .	60
Creating a new server cluster . . . . .	60
Modifying a server cluster . . . . .	61
Importing resource definitions . . . . .	62

**Chapter 7. Export, import, and promotion . . . . . 63**

Overview . . . . .	63
Job components that are migrated during export and import . . . . .	63
Exporting external references . . . . .	63
Export and import restrictions . . . . .	63
Recommended import order. . . . .	64
Security permissions when importing folders . . . . .	64
Exporting folders . . . . .	64
Importing folders . . . . .	65
Resolving import conflicts . . . . .	66
Promotion . . . . .	69
Promoting objects . . . . .	69
Promotion considerations. . . . .	69

**Chapter 8. Analytic data views . . . . . 71**

Creating analytic data views. . . . .	72
Data access plans . . . . .	73
Creating data access plans . . . . .	74
Adding mapped tables to analytic data views . . . . .	75
Modifying the data mapping for analytic data view tables . . . . .	77
Mapping table attributes to stream fields . . . . .	78
Overriding batch data access plan data sources . . . . .	78
Overriding real-time data access plan data sources . . . . .	79
Previewing data access plan data . . . . .	81
Deleting data access plans . . . . .	81
Data models . . . . .	82
Specifying the analytic data view editor advanced mode . . . . .	83
Adding unmapped tables to analytic data views . . . . .	84
Adding unmapped attributes to data model tables . . . . .	85
Defining relationships between analytic data view tables . . . . .	86
Adding derived attributes to analytic data view tables . . . . .	87
Modifying data model properties . . . . .	88
Deleting data model components . . . . .	90
Business object models . . . . .	90
Exporting data models as business object model archives . . . . .	91
Importing business object models as data models . . . . .	92
Execution object models . . . . .	93
Exporting data models as execution object model archives . . . . .	93

**Chapter 9. Scoring. . . . . 95**

Supported scoring functions and models . . . . .	96
Scoring configurations. . . . .	98
New scoring model configuration . . . . .	99

Model specific settings. . . . .	99
Data provider settings. . . . .	99
Input data order . . . . .	101
Input data returned settings . . . . .	101
Output data returned settings . . . . .	101
Logging settings . . . . .	102
Advanced settings. . . . .	104
Scoring configuration aliases . . . . .	105
Creating scoring configuration aliases . . . . .	105
Editing scoring configuration aliases. . . . .	106
Deleting scoring configuration aliases . . . . .	107
Scoring configuration association sets . . . . .	107
Creating scoring configuration association sets . . . . .	110
Editing scoring configuration association sets . . . . .	110
Deleting scoring configuration association sets . . . . .	111
A/B split testing for scoring configurations . . . . .	111
Evaluating scoring models by using A/B testing . . . . .	112
Scoring view . . . . .	113
Filtering the Scoring view . . . . .	114
Editing a scoring configuration . . . . .	114
Suspending and resuming scoring configurations . . . . .	114
Deleting a scoring configuration . . . . .	114
Scoring Graph view . . . . .	115

**Chapter 10. Jobs . . . . . 117**

What is a job? . . . . .	117
Versioning and labeling jobs . . . . .	117
Components of a job . . . . .	117
Prerequisite to running jobs . . . . .	118
External file dependencies . . . . .	118
Job process overview . . . . .	118
Working with jobs in the Content Explorer . . . . .	119
Creating new jobs . . . . .	119
Opening an existing job . . . . .	120
Viewing job properties . . . . .	120
Working with the job editor . . . . .	120
General information tab . . . . .	121
Job variables . . . . .	121
Adding steps to a job. . . . .	122
Executing job steps concurrently . . . . .	123
Specifying relationships in a job . . . . .	123
Sequential connector . . . . .	124
Pass connector . . . . .	124
Fail connector . . . . .	124
Conditional connector . . . . .	124
Viewing general relationship properties and modifying relationships . . . . .	125
Deleting relationships in a job . . . . .	126
Saving jobs . . . . .	126
Job step results . . . . .	127
Output file location . . . . .	127
Output file permissions . . . . .	127
Output file metadata . . . . .	128

**Chapter 11. Running jobs . . . . . 129**

On-demand job execution . . . . .	129
Specifying on-demand execution options . . . . .	129
Scheduled job execution. . . . .	130
Creating schedules . . . . .	130

Editing schedules . . . . .	133
Deleting schedules . . . . .	133
Message-based processing example . . . . .	133
<b>Chapter 12. Monitoring status . . . . .</b>	<b>135</b>
Accessing status views . . . . .	135
Selecting a server in status views . . . . .	135
Opening a job in the Job Editor . . . . .	136
Refreshing status views . . . . .	136
Reordering items in status views . . . . .	136
Deleting a job from status views . . . . .	136
Job schedule view . . . . .	136
Job history view . . . . .	137
Working with the job history table . . . . .	137
Model management views . . . . .	138
Model evaluation view . . . . .	139
Champion challenger view . . . . .	139
Predictors view . . . . .	140
Filters . . . . .	141
Filters common to all status views . . . . .	141
Job schedule filters . . . . .	142
Job History Filters . . . . .	143
Model management filters . . . . .	144
Server status view . . . . .	144
Setting the refresh rate . . . . .	145
<b>Chapter 13. Notifications and subscriptions . . . . .</b>	<b>147</b>
Notifications . . . . .	147
Job success and failure notifications . . . . .	147
Job step notifications . . . . .	148
Content notifications . . . . .	148
Notifications based on model evaluation return codes . . . . .	149
Label event notifications . . . . .	149
Notification settings . . . . .	150
Subscriptions . . . . .	154
Subscribing to files . . . . .	154
Modifying file subscriptions and unsubscribing from files . . . . .	154
Subscription email address . . . . .	155
Managing subscriptions . . . . .	155
Delivery failures . . . . .	155
<b>Chapter 14. Reporting job steps . . . . .</b>	<b>157</b>
General properties for reporting steps . . . . .	157
Data source for reporting steps . . . . .	157
Type for reporting steps . . . . .	158
Parameters for reporting steps . . . . .	159
Results for reporting steps . . . . .	159
Clean-up for reporting steps . . . . .	160
Notifications for reporting steps . . . . .	161
<b>Chapter 15. Visualization report job steps . . . . .</b>	<b>163</b>
General properties for visualization report steps . . . . .	163
Data source for visualization report steps . . . . .	163
Type for visualization report steps . . . . .	164
Parameters for visualization report steps . . . . .	164
Results for visualization report steps . . . . .	164

Clean-up for visualization report steps . . . . .	165
Notifications for visualization report steps . . . . .	166
<b>Chapter 16. SAS job steps . . . . .</b>	<b>167</b>
General properties for SAS steps . . . . .	167
Additional arguments for SAS steps . . . . .	168
Results for SAS steps . . . . .	168
SAS result formats . . . . .	169
Example SAS step . . . . .	169
<b>Chapter 17. General job steps . . . . .</b>	<b>173</b>
General properties for general job steps . . . . .	173
Input files for general job steps . . . . .	174
Input file source location . . . . .	174
Input file target location . . . . .	174
Output files for general job steps . . . . .	174
Output file intermediate location . . . . .	175
General job step examples . . . . .	175
Storing IBM Corp. PMML files . . . . .	175
Batch scoring using IBM SPSS Statistics PMML files . . . . .	176
Event-based scheduling . . . . .	178
<b>Chapter 18. Message-based job steps . . . . .</b>	<b>183</b>
General properties of message-based job steps . . . . .	183
<b>Chapter 19. Notification job steps . . . . .</b>	<b>185</b>
Adding a notification job step to a job . . . . .	185
General information . . . . .	186
Notification . . . . .	186
Updating notifications . . . . .	186
Selecting a new template . . . . .	186
<b>Chapter 20. Champion challenger job steps . . . . .</b>	<b>189</b>
Champion challenger overview . . . . .	189
Model evaluation metrics . . . . .	189
Order dependency . . . . .	189
General information . . . . .	190
Challengers . . . . .	191
Selecting challengers . . . . .	191
Invalid challengers . . . . .	192
Selecting challenger data sources . . . . .	192
Champion . . . . .	192
Testing the champion . . . . .	193
Data files . . . . .	194
Data view . . . . .	195
Enterprise view . . . . .	195
ODBC data sources . . . . .	195
Cognos import . . . . .	196
<b>Chapter 21. Submitted jobs . . . . .</b>	<b>197</b>
Restrictions within the Submitted Jobs folder . . . . .	197
Submitted jobs and expiration dates . . . . .	197
<b>Chapter 22. Accessibility features . . . . .</b>	<b>199</b>
Keyboard navigation . . . . .	199
Navigating the content explorer . . . . .	199
Navigating tables . . . . .	200

Navigating the job history and job schedule views . . . . .	200
Navigating the job editor . . . . .	200
Navigating the help system . . . . .	200
Accessibility for visually impaired users . . . . .	201
Accessibility issues for blind users . . . . .	201

**Chapter 23. IBM SPSS Collaboration and Deployment Services Enterprise View . . . . . 203**

IBM SPSS Collaboration and Deployment Services	
Enterprise View overview . . . . .	203
IBM SPSS Collaboration and Deployment Services Enterprise View builder interface . . . . .	203
Credential definitions . . . . .	204
Data source definitions . . . . .	204
Enterprise View . . . . .	204
Enterprise View editor . . . . .	204
Application View . . . . .	212
Creating an Application View . . . . .	212
Application View editor . . . . .	212
Merging view versions . . . . .	216
Resolve conflicts . . . . .	217
Data Provider Definition . . . . .	217
Creating a Data Provider Definition . . . . .	218
Data Provider Definition editor . . . . .	218
Data Provider Definition - Real Time . . . . .	221
Creating a Data Provider Definition - Real Time . . . . .	221
Data Provider Definition - Real Time editor . . . . .	222
Expression Editor . . . . .	227
Expression Panel . . . . .	227
Functions Panel . . . . .	228
Importing and exporting IBM SPSS Collaboration and Deployment Services Enterprise View objects . . . . .	253

IBM SPSS Collaboration and Deployment Services Enterprise View data types . . . . .	254
---	-----

<b>Notices . . . . . 261</b>	
Privacy policy considerations . . . . .	263
Trademarks . . . . .	263

**Glossary . . . . . 265**

A . . . . .	265
B . . . . .	265
C . . . . .	266
D . . . . .	266
E . . . . .	266
F . . . . .	267
G . . . . .	267
I . . . . .	267
J . . . . .	267
K . . . . .	268
L . . . . .	268
M . . . . .	268
N . . . . .	269
O . . . . .	269
P . . . . .	269
R . . . . .	269
S . . . . .	269
T . . . . .	270
U . . . . .	270
V . . . . .	270
W . . . . .	270
X . . . . .	271

**Index . . . . . 273**

---

## 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<sup>®</sup> 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.

---

## Chapter 2. What is new in this release

---

### What is new for IBM SPSS Collaboration and Deployment Services Deployment Manager users

IBM SPSS Collaboration and Deployment Services Deployment Manager 7 delivers new capabilities that can help you simplify deployment of predictive analytics and manage costs.

#### Scoring configuration aliases

A scoring configuration alias provides a fixed name to which scoring clients can send scoring requests. The alias passes the requests to an assigned scoring configuration. The scoring configuration that is assigned to the alias can be changed at any time, changing the model that is used for scoring.

#### Scoring configuration association sets

A scoring configuration association set consists of a primary scoring configuration and one or more alternative configurations for generating scores. Scoring requests that are submitted to the primary configuration are assigned to one of the configurations in the set according to specified distribution percentages.

#### Scoring configuration A/B split testing

Use A/B split testing to compare different scoring configurations in terms of overall performance. You can add the best configuration to your existing production tasks to improve their predictive capabilities.

#### Table caching for analytic data views

You can cache table values during data retrieval to prevent repeated retrieval of the same data.

#### Record limit for analytic data views

You can specify the maximum number of records to retrieve from a data source.



---

## Chapter 3. Getting started

---

### Launching IBM SPSS Collaboration and Deployment Services Deployment Manager

To launch the client:

1. From the Start menu, choose:

**All Programs > SPSS Inc. > IBM SPSS Collaboration and Deployment Services Deployment Manager**

The IBM SPSS Collaboration and Deployment Services Deployment Manager interface appears.

---

### Navigating through the system

IBM SPSS Collaboration and Deployment Services Deployment Manager relies primarily on tab-based navigation.

The interface is divided into the following major sections:

*Table 1. Sections in the IBM SPSS Collaboration and Deployment Services Deployment Manager.*

Section	Description	Location
Content Explorer	The Content Explorer is a tree that displays the contents of your repository. See the topic “Content Explorer overview” on page 13 for more information.	Left-hand-pane
Properties window	The Properties window displays the properties of the file that you select in the Content Explorer. Note that this Properties window is different from the individual properties that may appear as part of the Status tabs.	Lower-left-hand pane
Job Editor	The Job Editor allows you to drag-and-drop elements to create jobs.	Upper-right-hand pane

### Using the mouse 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.

### Dragging and dropping items in the user interface

You can drag items in the user interface. For example, you can reorganize items within the Content Explorer, or you can drag files from the Content Explorer into the Job Editor.

The system adheres to the following guidelines for drag-and-drop behavior.

- The Content Repository root cannot be moved.
- You can move items from the Content Explorer into the Job Editor. However, you cannot drag items from the Job Editor to the Content Repository. You must work with items within the Job Editor. See the topic “What is a job?” on page 117 for more information.

### Accessing help

Help is available via an online help system.

You can access help using any of the following methods:

**Help menu.** From the Help menu, choose **IBM SPSS Collaboration and Deployment Services Deployment Manager Help**.

**Dialog-level help.** To view online help in dialog boxes, click the **Help** button.

**F1 help.** In certain sections of the system context-sensitive help is available. To access context-sensitive help, press **F1**.

## Entry field content assist

The content assist feature provides job variables and predefined system property variables that can be used to insert values into entry fields. Variables available through content assist vary from field to field and may include timestamps, object paths, object URLs, job/step start and end times, execution identifiers for jobs and job steps, completion codes, variables passed to a job step by another job step, and variables defined at the job level. Entry fields that allow content assist are marked with a light bulb icon.

To insert variable values into a field, type \$. The list of available variables appears in a drop-down list. Click the variable name to display its description. Double-click the variable name to select it.

Date and timestamp variables allow the user to select the display format. To select a format after you have specified a variable, type . (period) after the variable name. The list of available formats appears in a drop-down list. Click the format name to display its description. Double-click the format name to select it.

Variables can be used to define file paths. However, a single backslash is an escape character for entry fields. Consequently, use a double backslash or a forward slash when specifying paths. For example, a path of

```
 ${JobVariable.Var1}\${JobVariable.Var2}
```

should be specified as

```
 ${JobVariable.Var1}\\${JobVariable.Var2}
```

or

```
 ${JobVariable.Var1}/${JobVariable.Var2}
```

## Field conventions

In IBM SPSS Collaboration and Deployment Services Deployment Manager, when property values of multiple objects or object versions are compared, property fields that have varying values are empty and marked by alert sign icons.

---

## Naming conventions within the system

At various points in the system, you are prompted to name items. For example, you will be asked to specify names for folders and jobs. All names within the system must be unique.

*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



---

## Exiting the system

You can exit the IBM SPSS Collaboration and Deployment Services Deployment Manager using either of the following methods:

- From the File menu, choose **Exit**.
- Click the close button (X) in the title bar of the user interface.



---

## Chapter 4. The Content Explorer

---

### Content Explorer overview

The Content Explorer is your entry point into the IBM SPSS Collaboration and Deployment Services Repository.

The Content Explorer contains a tree hierarchy of objects in the repository. The objects that appear in the Content Explorer depend upon your permissions. For example, you can view only folders for which you have permissions. In the Content Explorer, you can perform the following tasks:

- Log in to and log off from a server
- Establish server and credential definitions
- View object properties
- Access and work with files in the Content Repository

### What Is a content object?

A content object is any item that resides in the IBM SPSS Collaboration and Deployment Services Repository. Objects in the repository are stored in a relational database, such as Microsoft SQL Server or Oracle, in binary form.

You can add almost any type of file to the Content Repository. Examples of content objects include:

- IBM SPSS Modeler streams
- IBM SPSS Statistics syntax files
- SAS syntax files
- Jobs
- BIRT Report Designer for IBM SPSS reports

Typically, you perform an action on a content object. For example, you may add a IBM SPSS Modeler stream to a job, or you may run a job. At a minimum, you can move, copy, and paste content objects within the Content Repository.

### Organization of the Content Explorer

The server that hosts your IBM SPSS Collaboration and Deployment Services Repository is represented by a folder in the Content Explorer.

Within every server folder, the following items appear:

**Content Repository.** Contains all of your content objects. You can create subfolders within this folder.

**Submitted Jobs.** Displays the results of running reports using IBM SPSS Collaboration and Deployment Services Deployment Portal.

**Enterprise View.** A structure that provides a single, consistent view of enterprise data.

**Resource Definitions.** This folder contains server, credential, and data source definitions.

These items persist permanently at the root. You cannot move, copy, or delete these folders.

---

## Working with servers

When working with servers, you can perform the following tasks in the Content Explorer:

- Create new server connections
- Log in to a server
- Log off from a server
- Delete a server

### Creating new Content Server connections

Before you begin working, you must establish a connection to the server that houses your repository.

You need to create the server connection only once. After you have created the connection, the server folder appears in the Content Explorer, and you can simply log in to the server. See the topic “Logging in to a server” on page 15 for more information.

To create a server connection:

1. Launch IBM SPSS Collaboration and Deployment Services Deployment Manager .
2. From the File menu, choose:

**New > Content Server Connection**

The Create New Content Server Connection dialog box opens.

3. In the **Connection name** field, enter the name of your IBM SPSS Collaboration and Deployment Services Repository. This name is used in the Content Explorer at the root level.

*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

4. In the **Server URL** field, enter the complete connection URL for the server.

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 2. Example URL specifications.* This table lists some example URL specifications for server connections.

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

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

5. Click **Finish**.

## Logging in to a server

Before you begin working with the Content Explorer, you must log in to a server. The version of the content repository server must be equivalent to or greater than the IBM SPSS Collaboration and Deployment Services Deployment Manager version.

At a minimum, you must have one server connection defined in the Content Explorer. To log in to a server:

1. Double-click the server name. (Alternatively, you can expand the server folder by clicking on the + sign.) The Login to IBM SPSS Collaboration and Deployment Services Repository dialog box opens.
2. In the **User ID** field, enter a valid username for the server. Note that the username that you use determines what you see in the Content Explorer because different users may have different permission levels. See the topic “Modifying permissions” on page 28 for more information.
3. In the Password field, enter the password that corresponds to the username.
4. If multiple security providers have been configured for the server, use the Provider field to select the security provider against which to validate the user/password combination.
5. Click **OK**. The system opens the server folder and displays the expanded contents in the server directory.

Note: If single sign-on has been configured by the IBM SPSS Collaboration and Deployment Services administrator, the login page is bypassed and the repository is accessed without the user providing credentials. In that case, the user's Windows credentials are authenticated against an external directory service, for example, Windows Active Directory, which acts as a security provider for the system.

## Logging off from a server

You can access multiple servers simultaneously; however, before logging off from any of them, be sure to save changes to files or jobs on each server because the system does not prompt you to do so.

To log off from a server:

1. In the Content Explorer, right-click the server from which you want to log off.
2. Select **Logoff**. You are then logged off from the server, and the contents of the server folder are collapsed.
3. To access stored information, you must log back in to the server. See the topic “Logging in to a server” for more information.

## Changing server passwords

Username and password information is required to log in to a server in the Content Explorer. The password can be changed at any time. However, availability of the change password option depends on the security provider associated with the credentials.

For example, passwords can be changed when using IBM SPSS Collaboration and Deployment Services native security or IBM i providers, but not when using Active Directory.

The password modification takes effect immediately. Logging off and logging back in to the system is not required.

To change a server password:

1. Right-click the server name and select **Change Password**. The Change Password dialog box opens. This dialog box displays the name of the server to which you are currently logged in and your username. Neither the server name nor the username can be modified in this dialog box.

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, enter your new password again. The **OK** button is disabled until all of the fields in the dialog box contain data. In addition, the button is disabled until the information in the New Password and Confirm New Password fields is identical.
5. Click **OK**. The Change Password dialog box opens showing the message **Password changed successfully**.
6. Click **OK**.

---

## Working with files

In the Content Explorer, you can perform the following tasks:

- Open external files
- Add files to the repository
- Download files from the repository

### Opening external files

From the Content Explorer, you can open and view files.

The method used to open and display the file depends on the file type. For example, if you open a text file, the text is displayed in the Job Editor. However, if you double-click a IBM SPSS Modeler stream, the system launches the IBM SPSS Modeler application.

### Permissions inheritance

It is important to understand the relationship between IBM SPSS Collaboration and Deployment Services Repository resources, particularly when you copy or move resources.

The following guidelines apply:

- **Creating resources.** When a resource is added to the repository (for example, when a new job is created), it inherits the permissions of its parent folder; a user with *Write* permissions to a folder he will by default have the same permissions to any resources created in it.
- **Copying resources.** When you copy a resource to a new folder, the resource retains the permissions of the original folder. However, if the user copying the resource is not the owner of the original resource, then the system changes the ownership of the resource to the new user.
- **Moving resources.** When you move a resource from one folder to another, the resource retains the permissions of the original folder. When you cut and paste a resource, this is considered a move. Thus, the resource retains the permissions of the original folder.

---

## Working with the repository

Files can be added to the repository or downloaded from the repository. Typically, individual files are added or downloaded.

It is important to note that adding or downloading files is different from importing or exporting files. See the topic "Overview" on page 63 for more information.

### Adding files to the repository

You can add almost any type of file to the repository.

**Using the File menu.** To add a file to the repository:

1. In the Content Explorer, select the folder to which you want to add the file.

2. From the File menu, select **Add File to Repository**. The Add Files dialog box opens. If the **Add File to Repository** option is disabled, then you have clicked on an object in the Content Explorer and not on a folder.
3. Navigate to the file that you want to add to the repository.
4. Click **Open**.

**Dragging and dropping files.** Alternatively, you can drag files into the repository.

## Downloading files from the repository

Using the Content Explorer, you can download files from the repository to another computer.

To download files:

1. In the Content Explorer, select the file that you want to download.
2. From the File menu, choose **Download File**. If your file has multiple versions, the Select File Version dialog box opens. See the topic “Selecting a version” on page 31 for more information. Otherwise, the Download Files dialog box opens.
3. Navigate to the folder in which you want to place the file.
4. Click **OK**. A copy of the file is saved in the folder that you specified.

## Deleting files from the repository

Objects can be deleted from the repository individually or in bulk, provided the appropriate permissions are specified.

However, although multiple IBM SPSS Collaboration and Deployment Services Enterprise View objects can be selected and deleted simultaneously, IBM SPSS Collaboration and Deployment Services Enterprise View objects cannot be deleted in conjunction with non-IBM SPSS Collaboration and Deployment Services Enterprise View objects. If an IBM SPSS Collaboration and Deployment Services Enterprise View object is selected along with a non-IBM SPSS Collaboration and Deployment Services Enterprise View object, the delete process will fail. See the topic “IBM SPSS Collaboration and Deployment Services Enterprise View overview” on page 203 for more information.

---

## Searching

Contents of the IBM SPSS Collaboration and Deployment Services Repository are searchable. The results appear on the Search Results tab.

The following types of searches are available:

- Simple search.
- Advanced search.

It is important to note that the search function will search for full text matches only. Currently, partial text and wildcard searches are not supported. Finally, searches cannot be saved and do not persist.

## Guidelines for searching

When searching for objects in the system, the following guidelines apply:

- In order to find an object, the name provided must match the repository object's name exactly.
- The search feature looks inside of IBM SPSS Modeler streams to find nodes that contain the search string.
- Although the following objects are returned if they match a search string exactly, the search feature does not look *within* the following types of objects: scenario files, submitted jobs, and resource definitions.

- The following characters are not allowed in search strings:
  - Single quotation marks (')
  - Double quotation marks (")
  - Parentheses ( or )
- Certain terms are excluded from the search. See the topic “Terms Not Included in a Search” on page 21 for more information.

## Accessing the search dialogs

To access the search dialogs:

1. Ensure that you are logged in to the server that contains the repository that you want to search.
2. In the Content Explorer, select any server folder. The search must be conducted at the server instance level. The search option is disabled at the subfolder level.
3. From the Edit menu, choose **Search**. By default, the simple search dialog appears. If the *Search* option is not enabled, then the appropriate folder level was not selected. The search option is enabled only at the server level. All subfolders are automatically searched.
4. To perform an advanced search, click **Advanced**. See the topic “Advanced Search” for more information.

## Simple Search

A simple search looks through the repository for objects that match the specified string and returns all matching results. The simple search does not distinguish among property types (e.g., author or title) or narrow results by time frame.

To search for a string:

1. In the **Search** field, enter a text string. Quotation marks are not necessary.
2. Click **Search**. The Search Results tab is populated with all of the objects that meet the search criteria. See the topic “Viewing Search Results” on page 21 for more information.

## Advanced Search

The advanced search allows greater refinement of the search. For example, in addition to text strings, the advanced search option provides the ability to search by date range.

To perform an advanced search, the following options are available:

### Property

The list of searchable properties depends on the types of files contained within the repository and may include:

- Author
- Description
- Keyword
- Label
- MIME type
- Object last modified by
- Parent URI
- Title
- URI
- Version created by
- Version URI



For any property selected, an exact value needs to be specified. A list of options will not appear for any of these property types. For URI properties, the URI must use the id and version marker specification. Paths and version labels cannot be used when searching by URI.

To search by property:

1. From the **Property** drop-down list, select a property type.
2. In the **Value** field, type a value that corresponds to the property type selected. Quotation marks are not necessary.
3. Click **Add**. The property type and the value appear in the Search Terms box. For the initial entry in the Search Terms box, no AND or OR grouping is specified. By default, any subsequent search terms are combined with the AND operator. The AND operator can be changed to OR at any time. See the topic “Switching Between AND and OR” for more information.

## Date Search

A date range can also be specified in an advanced search. Valid parameters for a search by date range include:

**Date search.** Valid values include: *Expiration Date*, *Last Modified*, and *Version Created On*.

**Date range.** A date range must be provided. The date in the **To** field must be equal to or greater than the date in the **From** field.

**Time range.** The time range is optional. However, to search by a time range a date range must be specified first.

## Refining Search Terms

After parameters are specified in the property or date fields and added to the Search Terms list, the selections can be further refined.

Specifically, the following enhancements can be made:

- Switching between *AND* or *OR*.
- Grouping search terms
- Editing search term values.
- Reordering search terms.
- Deleting search terms.

### Switching Between AND and OR:

By default, search terms are joined by the Boolean operation *AND*. The Boolean operator can be changed from *AND* to *OR* and vice versa.

The button that is available depends upon the search term selected. If the search term is connected by *AND*, then the **OR** button appears. If the search term is connected by *OR*, then the **AND** button appears.

To change the connector:

1. Select the search term whose connector you want to change.
2. Click **AND** or **OR**.

### Grouping and Ungrouping Search Terms:

Search terms can be grouped to refine a search further. The ability to perform AND and OR searches allows searches to be organized across attributes. Grouping allows AND and OR searching to occur within the same attribute. Nested groups are supported.

For example, suppose that you want to search for an object modified in December 2007 that could have been created by one of two authors (Joe Author or Jane Author) and carries one of two labels (Test 1 or Test 2). By default, all search terms are joined by *AND*. A search for objects authored by Joe Author and Jane Author *AND* labeled Test 1 and Test 2 might return limited results. By changing the *AND* between authors and labels to *OR* and using grouping, the search could be refined so that the search results contain a list of objects authored by either Joe Author or Jane Author and labeled either Test 1 or Test 2 and last modified in December 2007.

Suppose the following property values were set and appeared in the Search Terms list:

```
'Author' = 'Joe Author'  
OR 'Author' = 'Jane Author'  
AND 'Label' = 'Test 1'  
OR 'Label' = 'Test 2'  
AND LastModified BETWEEN '12/1/07' AND '12/31/07'
```

When grouped the properties would be organized in the Search Terms list as follows:

```
('Author' = 'Joe Author'  
OR 'Author' = 'Jane Author')  
AND ('Label' = 'Test 1'  
OR 'Label' = 'Test 2')  
AND LastModified BETWEEN '12/1/07' AND '12/31/07'
```

To group search terms:

1. Select the search terms that should be evaluated as a group. To group selected items, the items must be adjacent to each other on the list of search terms. To select multiple terms, press **Ctrl** while selecting rows from the list.
2. Click **Group** or **Ungroup**. The *Ungroup* option only appears if the search terms selected were previously grouped together.

### Editing Search Terms:

The values for search terms can be edited. However, the property type cannot be edited. To add a new property type to the Search Terms list, a new property must be selected from the **Property** drop-down list and added to the Search Terms list.

To edit a previously specified property value:

1. Select the search term to edit.
2. Click **Edit**. The Edit Search Value dialog appears. The contents of the dialog may vary depending upon the search term selected. For example, the dialog for the *Author* property type contains a text field for the author's name; whereas, the dialog for the *Last Modified* property type contains a date range.
3. Specify a new value for the property type.
4. Click **OK**.

### Reordering Search Terms:

The order in which search terms appear on the list can be reordered. The order of search terms is only important when grouping terms together. Search terms must be adjacent to each other on the search terms list in order for grouping to be possible.

However, the order in which search terms (grouped or ungrouped) appear on the search terms list does not matter. The search feature will return results in order of relevancy. Thus, the search term that matches the largest number of repository objects will be listed first in the search results. The search term with the second largest number of hits will be listed second, and so forth.

To reorder items in the search terms list:

1. From the Search terms list, select the search term to move.

2. Click **Up** or **Down** until the search term appears in the target location. It is important to note that after rows are grouped, the movement of rows is limited. Movement that would potentially affect a grouping is not allowed.

### Deleting Search Terms:

Search terms can be deleted from the search terms list. After a search term is deleted from the list, it cannot be recovered. The search term would have to be re-created.

To delete a search term:

1. Select the search term to delete.
2. Click **Delete**. The search term is removed from the list.

## Viewing Search Results

By default, the Search Results table is hidden. When a search is executed, the Search Results table appears.

To open the Search Results table independently, choose **Search Results** from the View menu.

The search results table contains the following information.

*Table 3. Search results.*

Column	Description
Name	The name of the file that contains the search string.
Folder	The folder that contains the file.
Author	The user who created the file.
Type	The file type. External applications are prefaced with <code>application/x-vnd</code> . For example, the type for a IBM SPSS Modeler stream is <code>application/x-vnd.spss-clementine-stream</code> .
Last Modified	Date and time of the most recent modification to the returned item.
Version Count	The number of versions of the search object.

## Terms Not Included in a Search

Stopwords are excluded from searches. The words in the following alphabetical list are not indexed:

a, all, am, an, and, any, are, as, at, be, but, by, can, could, did, do, does, etc, for, from, goes, got, had, has, have, he, her, him, his, how, if, in, is, it, let, me, more, much, must, my, nor, not, now, of, off, on, or, our, own, see, set, shall, she, should, so, some, than, that, the, them, then, there, these, this, those, though, to, too, us, was, way, we, what, when, where, which, who, why, will, would, yes, yet, you

As a result, the system ignores them when searching for items. Currently, you cannot modify the terms in this list.

---

## Object Locking

In the Content Explorer, objects can be locked to prevent other users from making changes.

For example, suppose that a file associated with a step within a job needs to be modified. Locking the file would prevent others users from making changes that could potentially conflict with the current user's changes. If another user opens a locked object, the new user receives a message indicating that the file is locked and can only be opened in read-only mode.

The following guidelines apply to locked objects:

- Most objects in the repository--including resource definitions--can be locked. However, folders cannot be locked. In addition, objects within the *Submitted Jobs* folder cannot be locked.
- When an object is locked, all versions of that object are locked. Thus, properties of that object cannot be edited. However, properties will be visible in read-only mode.
- Locks, which are denoted by the lock icon, persist across sessions.
- Objects remain locked until the lock is explicitly removed. Expiration dates cannot be associated with a lock.
- A locked object cannot be renamed, moved, or deleted. However, a locked object can be copied. The lock is not copied along with the object.
- Locks do not persist from one repository instance to the next. If a locked object is exported, the lock is not exported along with the object.
- Any user can lock an object. Users can also unlock their own objects. However, to unlock additional objects, users must have the action, *Manage Locks*.

## Locking Objects

To lock an object:

1. In the Content Explorer, right-click the object and select **Lock**. The lock icon appears for the object.

## Viewing the Locked Objects Table

The Locked Objects table contains a list of locked objects and corresponding details about those objects.

Specifically, the Locked Objects table contains the following information:

**Name.** The name and path of the object that is locked.

**Locked By.** The user who locked the object.

**Locked Since.** The date and time at which the object was initially locked.

To access the Locked Objects table:

1. From the View menu, select:

**Show View > Locked Objects**

## Unlocking Objects

Objects can be unlocked using either of the following methods:

**Content Explorer.** To unlock an object, right-click the object and select **Unlock**. The object is unlocked, and the lock icon is removed.

**Locked objects table.** To unlock an object:

1. Navigate to the Locked Objects table.
2. Select the object(s) to unlock. To select multiple rows, press the **Ctrl** key and click simultaneously.  
*Note:* Only objects for which the user has appropriate permissions can be unlocked.
3. Click the **Unlock** icon. The selected objects are unlocked and removed from the Locked Objects table.

---

## Obtaining updates

Periodically, updates may be available on the server. Each time IBM SPSS Collaboration and Deployment Services Deployment Manager is launched, the system checks the server for updates.

If updates are available, the system automatically applies the updates and sends a message indicating that the updates have been applied. After updates have been applied, the IBM SPSS Collaboration and Deployment Services Deployment Manager needs to be restarted for the changes to take effect.



---

## Chapter 5. Properties

---

### Working with object properties

Properties are metadata associated with a content object, which describe the overall content object. Author and modification date are characteristics associated with a content object. Typically, more properties are available as you delve deeper into the content tree.

For example, for a folder at the root of the content tree, *title* is the only property available. However, when you view properties for a IBM SPSS Modeler stream, you see additional properties, such as author, description, version, and so on.

When working with properties, you can:

- View properties
- Edit properties

### Viewing object properties

When you select an object in the Content Explorer, properties appear in the Properties pane.

The type and number of properties that are displayed depend upon the object and your location within the content tree. For example, if you click a server folder at the root level, you will see a title in the Properties pane.

Although you view object properties in the Properties pane, you edit them in the Properties dialog box.

*Table 4. Object properties.*

Property	Description	Can you view this in the Properties pane?	Can you modify this in the Properties dialog box?
Author	Author refers to the creator of the job. Typically, the author corresponds to a user ID in the system and not an individual person. For example, the author of the job may be the administrator.	Yes	No
Description	The description displayed here corresponds to the description that you specified when you created the object.	Yes	Yes
Modification date	The modification date and time of the most recent modification to the object is displayed here.	Yes	No
Title	The title is the name of the object.	Yes	Yes. Title corresponds to name in the Properties dialog box.
Version	Identifier that distinguishes one version of an object from another. The identifier consists of the version number, a colon, the version modification date, and the version modification time.	Yes	Yes
Name	The name of your object.	Yes. Name corresponds to title in the Properties pane.	Yes

Table 4. Object properties (continued).

Property	Description	Can you view this in the Properties pane?	Can you modify this in the Properties dialog box?
File type	The MIME type associated with the job.	No	Yes
Last modified by	The last user who made changes to the job.	No	No
Custom properties	Properties defined in the Custom Properties dialog box. These properties will vary as they are specific to each server instance.	No	Yes

## Editing object properties

You may want to modify properties. For example, you may want to add keywords, which are used by the system to search for objects.

Properties are grouped into the following categories:

- General
- Labels
- Permissions
- Versions

To access the Properties dialog box:

1. In the Content Explorer, right-click the object whose properties you want to view.
2. Choose **Properties**. The Properties dialog box opens.

### Editing general properties

To edit general object properties:

1. In the Content Explorer, right-click the object whose properties you want to view.
2. Choose **Properties**. The Properties dialog box opens. By default, the General properties dialog box opens.

The General properties dialog contains information about the object that you selected. Specifically, this dialog describes the following information, which is modifiable, unless otherwise noted.

**Title.** The name of the object.

**Object URI.** The uniform resource identifier (URI) for the object. This value is generated by the system and cannot be modified.

**Author.** The user ID under which the object was created.

**Content type.** The object type. A new object type can be specified for most objects in the repository. However, a new content type cannot be applied to the following objects:

- A job
- Any object in the *Resource Definitions* folder
- Any object in the *Submitted Jobs* folder

*Note:* Properties can be modified for an object in the *Submitted Jobs* folder once the object has been moved out of that folder into the Content Repository. See the topic Chapter 21, “Submitted jobs,” on page 197 for more information.



To change the file type that is associated with the object:

1. Click the ellipsis button. The File Types dialog box opens.
2. From the Type list, select the new file type.
3. Click **OK**. The new file type appears in the General properties dialog box.

**Last modified by.** The user ID under which the last modification was made. This value is generated by the system and cannot be modified.

**Modification date.** The date and time at which the last modification occurred. This value is generated by the system and cannot be modified.

**Custom properties.** If custom properties were defined for this server instance, custom property values can be assigned in this dialog box.

**Topics.** If topics were defined for this server instance, topic values can be assigned in this dialog box.

#### **Setting values for custom properties:**

After a custom property has been created, you can assign values to the custom properties fields. Custom properties appear in the General properties dialog box.

It is important to note that you may not see custom properties for all objects on a server. Custom properties can be restricted to certain objects (for example, jobs) or certain file types.

To set values for an existing custom property:

1. Right-click on any object within the server and choose **Properties**. The General properties dialog box opens.
2. Specify the values for the custom properties that appear. The options that appear depend upon the types of custom properties that were created. See the topic “Creating custom properties” on page 39 for more information.
3. To save your changes, click **OK**.

#### **Assigning topics:**

After you have created and saved a topic definition, you can assign the topic to a content object.

To assign a topic:

1. Navigate to the General Properties dialog box. The **Topics** table lists any previously assigned topics.
2. To assign a new topic, click **Add**. The Add Topics dialog box opens. This dialog box displays the topic hierarchy that was established in the Topics Definition dialog box.
3. Select a topic from the hierarchy.
4. Click **OK**. The Properties dialog box reopens, and the topic is added to the **Topics** table. Note the format of the topic reference. Each slash in the topic reference represents a folder in the topic hierarchy. The first slash represents the uppermost topics folder.

*Removing topics:* To remove a topic assignment:

1. Navigate to the General Properties dialog box. The **Topics** table lists any previously assigned topics.
2. Select a topic from the **Topics** table.
3. To remove a topic, click **Remove**. The topic is removed from the table.

## Modifying permissions

Each object has a permission level associated with it. At times, you may want to modify permissions. For example, if you have added steps to a job and scheduled its execution for the next six months, you may want to restrict other users to read-only permissions to prevent changes to the job.

To access the permissions for the object:

1. In the Properties dialog box, click the **Permissions** tab. A list of users and their corresponding permissions appears.

You can perform the following tasks in this dialog box:

- Add a new user or group
- Modify permissions for an existing user or group
- Delete an existing user or group

**Adding new users or groups:** To add a new user or group and assign permissions:

1. Click **Add**. The Select User or Group dialog box opens.
2. From the **Select provider** drop-down list, select the entity that contains your user and group information. The default is *Local User Repository*.
3. In the Find field, type the first few letters of the user ID you want to add. To search for all available user IDs, leave this field blank.
4. Click **Search**. The list of users and groups that correspond to your search appears in the dialog box.
5. Select a user or group from the list.
6. Click **OK**. The user or group appears in the Principal list in the Permissions dialog box.

**Modifying permissions for an existing user or group:** To modify permissions:

1. Click in the Permissions cell that corresponds to the user or group whose permissions you want to modify.
2. From the drop-down list, select a new permissions level. You can select from the following options:
  - Read
  - Write
  - Delete
  - Modify Permissions

### Cascading permissions:

After permissions have been established in the Permissions table, the permissions changes can be cascaded, which means that the content of the current folder will have the same permissions as the parent folder. Permissions can be cascaded completely, partially, or not at all. By default, permissions are not cascaded.

*Note:* The ability to cascade permissions to an object is controlled by the Modify Permissions option. Permissions are cascaded to expired objects as well as active objects. If an object is not visible because it is expired and the user is not assigned the View Expired action but has Modify Permissions permission, permissions will still be cascaded. If the user does not have Modify Permissions permission, then cascading permissions will not be applied.

To cascade permissions:

1. Select the level of cascading from the drop-down list. Valid values include:
  - Do not cascade permission changes to child folders and content.
  - Cascade only the updated permissions to all child folders and content.
  - Cascade all permissions to all child folders and content.

2. Click **OK**.

**Deleting a user or group from the permissions list:** To delete a user or group:

1. Click in the Principal cell of the user or group that you want to remove.
2. Click **Delete**. The Delete Confirmation dialog box opens.
3. Click **OK**. The user is removed from the Permissions list.

## Working with labels

Version labels allow a version of an object to be distinguished from another version of the same object, using a user-defined name instead of a system-generated identifier. You can use a label to reference the specific version without having to determine the version identifier.

For example, version labels are useful when scheduling jobs. Associating the schedule with a specific label ensures that only the version of the job with that label is executed when the schedule activates. Any other versions of the job, including those that were created after the labeled version, are ignored by the schedule.

Although any particular object version can be associated with multiple different labels, labels must be unique across the versions of an object. For example, there cannot be two versions of a particular job that are both labeled *Production*. This restriction allows a one-to-one mapping of label references to object versions.

Labels can be moved from one version to another version. When you move a label, every item that references the object with that label will use the new version of the object instead of the old. For example, if you move the *Production* label of a report from the first version to the seventh, any job that references the *Production* version of the report automatically uses the seventh version instead of the first.

The system automatically associates the internal label *LATEST* with the most recent version of an object. Any time a new version of an object is created, the *LATEST* label is moved to the new version automatically. Because this behavior is automatic, referencing a version by using the *LATEST* label is discouraged. Doing so often results in undesirable results. For example, if a job references the *LATEST* version of a report, the most recent version is executed. The most recent version may be a working version that is not yet finalized. However, if the job references the *Production* version, the job ignores every interim version created after the *Production* version. When a new version is ready for general use, the *Production* label can be moved and the job will execute that version.

Version labels persist until they are explicitly changed or removed. It is important to note that removing a version label does not remove the label from the repository. Removing a version label simply severs the connection between an object and the version label. The label is still available for use by other objects.

The following information appears on the Labels tab of the Properties dialog box:

- **Version.** The version identifier, which consists of the version number, suffixed with the date and time that the version was created. The version identifier cannot be modified.
- **Label.** The labels associated with the version. If a version has multiple labels, the labels appear in this column separated by commas. If the column width prevents all of the labels from being displayed, position the mouse cursor over the Labels cell to display the complete list. Apply and remove labels by using the Edit Version Labels dialog box. See the topic “Editing the version label” on page 30 for more information.
- **Creation Date.** The date and time on which the version was created. The creation date cannot be modified.
- **Expiration Date.** The date on which the version is set to expire. The expiration date is set on the Versions tab of the Properties dialog box. See the topic “Version properties” on page 37 for more information.
- **Created By.** The user who created the version. This value cannot be modified.

From the Labels tab, you can perform the following actions:

- Apply a new version label
- Edit an existing version label
- Apply a predefined version label
- Specify multiple version labels
- Remove a version label
- Delete a version

#### **Version label recommendations:**

A label policy that defines a set of labels and outlines how they should be used is strongly recommended. Such a policy will help prevent inadvertently breaking automated processes as objects change.

In particular, this policy should consider the following:

- When creating new objects and object versions, apply specific version labels, such as *Test* or *Production*, that associate the version with a procedural stage. The label should be moved to a new version when that version is ready for general use. This ensures that references to the object will consistently be linked to the expected version and the references will not change without an explicit label move.
- Job versions should be identified with specific labels that allow execution of an intermediate test version of the job while the validated production version continues to operate. This allows new objects or object versions to be tested before being promoted to general use. The test version of the job can be promoted by moving the version label for production when the test job is ready for general use.
- Job steps should always reference object versions by using a specific version label instead of the system *LATEST* label to ensure that a specific object version will be used regardless of whether more recent versions of that object exist. This permits the job to run successfully while the referenced object is being revised. If the reference uses the *LATEST* version, the job will use the most recent object version, which may require modifications of the job properties. Without those changes, the job could fail. By referencing a specific version label and by labeling job versions, you can modify the job properties to accommodate changes to the referenced object version. Once those changes are validated, you can promote the new object version and the revised job for production use.
- Use label security, which allows for restricting the users allowed to assign or move specific version labels, to ensure that a proper review and validation of objects and jobs by a qualified individual occurs before new versions are promoted to production use.

#### **Editing the version label:**

To edit the version label, select the object version from the table on the Labels tab of the Properties dialog box and click **Edit**.

To apply a new label, perform the following steps:

1. Type the name in the **New Label** field. Labels cannot contain commas, apostrophes, semicolons, double quotation marks, greater than symbols, or less than symbols.
2. Click the arrow button to add the label to the **Applied Labels** field.

To apply an existing label, perform the following steps:

1. Select the label(s) from the **Available Labels** list. To select multiple labels, press **Ctrl** and make multiple selections.
2. Click the arrow button to add the selected label(s) to the **Applied Labels** list.

To remove the association between a version and a label, perform the following steps:

1. Select the label from the **Applied Labels** list.
2. Click the arrow button to remove the label from the **Applied Labels** list.

When your modifications are complete, click **OK**. The Versions table displays the updated values.

#### *Resolving duplicate version labels:*

Although multiple version labels can be applied to an object, each individual version label must be unique for that object. (However, a label with the same name can be used for another object.) In the event of a conflict, the Labels Already Exist dialog box describes the version label duplication.

To resolve a duplicate label conflict, the following options are available:

- **Apply the label.** To apply the label to the existing version and remove the label from the previous version, click **OK**.
- **Specify a new label.** To maintain the label with the previous version and specify a new label for the current version, click **Cancel**. The Edit Version Labels dialog reappears. Specify a new label in this dialog.

**Deleting versions:** To delete an IBM SPSS Collaboration and Deployment Services Repository object version:

1. On the Labels tab of the Properties dialog box, select the version.
2. Click **Delete**. The Delete Confirmation dialog box opens.
3. Click **OK**. The version is removed from the list.

*Note:* Once a version is deleted from IBM SPSS Collaboration and Deployment Services Repository, it cannot be recovered.

#### **Selecting a version:**

At various points in the system, you may be asked to select a version of a file. For example, if you want to download a file that has multiple versions, the system prompts you to select a version.

To select a version:

1. If your file has multiple versions, the Select file version dialog box opens.
2. Select the version of the file on which you want to perform an action.
3. Click **OK**.

If the selected version of the file has multiple labels associated with it and the version opens in the IBM SPSS Collaboration and Deployment Services Deployment Manager environment, the first label appears with the name of the file in the tab title to assist in identifying the version.

## **Working With Expiration Dates and Expired Files**

### **Working with Expiration Dates and Expired Files**

An expiration date specifies the date after which a file is no longer actively in use. When an item expires, it is no longer visible to the general user. However, the file owner and administrators can see expired files.

Note that expiration does not equal deletion. An expired file is not removed from the repository. By default, files do not expire. An expiration date must be explicitly set. Expiration dates can be set, modified, or reactivated by any of the following:

- The owner of the file
- An administrator
- Any user with write permissions to the file

The file or job itself does not have an expiration date associated with it. Expiration dates are set at the version level. Version numbering continues sequentially. However, a gap in version numbers may appear on the versions list if some file versions have expired.

When a file version expires, any corresponding labels are unaffected. If an expired file version is opened, an error message appears, indicating that the file has expired. When the labeled version of a job expires, it cannot be executed after the expiration date. Thus, any schedules that contain the expired job version are affected.

Schedules are not deleted when the associated version expires. The system will still attempt to execute the scheduled job at the designated time. However, because the job version contained in the schedule has expired, the job will fail, generating an error message in the log. The failure of the scheduled job is recorded in the job history. See the topic “Job history view” on page 137 for more information.

### Setting and Modifying Expiration Dates:

Expiration dates are set in the Properties dialog box. In addition, an expiration date can also be specified when a job is initially created.

By default, objects do not expire. Expiration dates can be modified at any time. In addition, expiration dates can be retroactive. To set an expiration date on an object:

1. In the Content Explorer, right-click on the object and select **Properties**. The Properties dialog box opens.
2. Click the Versions tab and select the version on which to specify an expiration date.
3. Click the ellipsis button next to the **Expiration Date** field. The Expiration Date dialog box opens. By default, no expiration date is set.
4. Select one of the following options:
  - **No Expiration Date**. The version does not expire.
  - **Expire in**. The version expires at a time relative to today's date—for example, in 30 days. Valid relative time periods include days, weeks, months, or years. The maximum expiration date allowed is 10 years.
  - **Expire on**. The version expires on the date selected.
5. Click **OK**. The expiration date appears in the **Expiration Date** field. For both relative and specific dates, the exact calendar expiration date appears in this field. If no expiration date was specified, the field remains blank.

Alternatively, the expiration date can be specified when a new job is created. The process for accessing and supplying information for the expiration date dialog box is the same. However, the dialog box is accessed from the New Job Information dialog box instead of the Properties dialog box.

### *Importing and Exporting Expiration Dates:*

When files are imported into and exported out of the system, the expiration date attached to a file version is included as well.

### *Restrictions on Expiration Dates:*

Expiration dates can be applied to most files within the Content Explorer. However, expiration dates cannot be applied to objects in the Resource Definitions folder—for example, server, credential, and data source definitions.

### Viewing Expired Files:

Expired files are visible depending upon whether or not all versions of a file have expired.

**Files with some versions expired.** For files with some expired versions, the expired versions are displayed.

**Files with all versions expired.** The expired files are hidden in the Content Explorer. However, if the owner of the expired files or an administrator views the tree, the expired files are visible.

*Searching for Expired Files:*

The process of searching for expired files is the same as the process of searching for other objects in the Content Explorer.

See the topic “Searching” on page 17 for more information.

To see expired files in the search results, a user must be one of the following:

- A user with the *view expired files* action
- The owner of the file
- An administrator

**Reactivating Expired Files:** Expired files can be reactivated. To reactivate an expired file version:

1. Navigate to the existing expiration date in the Properties dialog box.
2. Modify the expiration date to a future date.

---

## Working with server properties and user preferences

Server properties are divided into the server properties and user preferences subcategories.

Both types of properties apply to the current server instance (or user) only. To access the Properties dialog box for a server instance, right-click on a server name in the Content Explorer and choose **Properties**. The Server Properties dialog box opens.

### Server properties

Server properties are similar to content object properties. However, server properties apply to the overall server instance.

Users with the appropriate access privileges can view and modify the following:

- Custom properties. See the topic “Working with custom properties” on page 38 for more information.
- Server connection
- Label security
- Topic definitions. See the topic “Working with topic definitions” on page 42 for more information.

### Server connection

Server connection properties define the information that is needed to connect to a IBM SPSS Collaboration and Deployment Services Repository. To access the connection properties, click the + to expand the server properties options in the Server Properties dialog box and select Server Connection.

**Name** The name for the server connection.

#### Server URL

The complete connection URL for the server that hosts the IBM SPSS Collaboration and Deployment Services Repository.

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

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

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

## Label security

Label security provides control over which principals can view or modify user-defined version labels in the system, allowing the specification of which users can move or delete a label. Although users may be able to view a version of a resource in the IBM SPSS Collaboration and Deployment Services Repository, only those users with permissions for the labels associated with the version can access the labels.

Moreover, label security applies to any version of any resource using the label. For example, to control who can assign production versions of resources, an administrator may restrict access to the *Production* label to prevent other users from moving that label from one version to a newer version. The security defined for the *Production* label applies to every version of every resource using that label.

To access label security properties, click the + to expand the server properties options in the Server Properties dialog box and select Label Security.

The Labels table lists all labels in the system you can access. Select a label to view the associated principals and their permissions

For a selected label, the Principals table identifies label permissions for principals in the system. A principal can have one of the following permissions for a label:

- **Use versions.** Any principal with this permission can view the version of a resource associated with the label by referencing the label.
- **Manage label.** Principals with this permission can apply, move, and delete the label, as well as delete the version of a resource having the label. If a user does not have the *Manage label* permission for a label, that user cannot delete a resource or version of a resource to which the label is applied

To edit the security for a label, such as adding or deleting principals, you must have the *Manage label* permission for the label. Administrators have this permission for all labels in the system. The *Manage label* permission does allow you to alter your own permission for a label, making it possible for you to edit your permissions for a label in such a way that you will be unable to work with the label in the future. For example, you might change your permission to *Use versions*, which would subsequently prevent you from changing the security for a label. If situations such as this occur, contact your administrator to have your permissions adjusted as needed.

## Adding principals for a label

To add a principal for a label:

1. Select the label from the Labels table.



2. Click **Add**. The Select User or Group dialog box appears.
3. Select the principal to add.
4. Click **OK**.

The principal appears in the Principals table with the default permission.

### Modifying permissions for principals

To modify the label permissions for a principal:

1. In the Principals table, click the Permissions cell for the principal.
2. Select the permission.

### Deleting principals for a label

To delete a principal for a label:

1. Select the label from the Labels table.
2. In the Principals table, select the principal to be deleted. To select multiple principals, hold down the Ctrl key while selecting principals.
3. Click **Delete**.

The principal is removed from the Principals table and will be unable to access the label.

## User preferences

User preferences apply to the user ID that is currently logged in to the server.

Users with the appropriate access privileges can view and modify the following:

- Analytic data view editor advanced mode
- Default permissions
- Distribution channels
- Subscriptions recipient
- Version label

### Default permissions

The default permissions option allows a user to set up a default permissions for new files and folders that are created. Thus, when a user creates a new object, such as a job, the object's permissions list defaults to the list of principals and permissions defined in the user's preferences for the instance of IBM SPSS Collaboration and Deployment Services.

In the Default Permissions dialog box, the following changes can be made:

- Add a new principal or delete an existing principal
- Modify permissions for a principal

To specify default permissions:

1. Navigate to the Properties dialog box for the server instance.
2. In the Properties tree, expand the User Preferences list and select **Default Permissions**.
3. Modify the list of principals as needed. Users or groups can be added to or deleted from the Principal list.

**Add a new principal.** To add a new principal, click **Add**. The Select User or Group dialog box opens. Select a user or group from the list and click **OK**.

**Delete an existing principal.** To delete an existing principal, select a principal from the list and click **Delete**. The Confirm Deletion dialog box opens. Click **OK**.

4. To modify the permissions associated with a principal, click the drop-down arrow in the Permissions column and select the permission level to associate with the principal. Valid values include:
  - Read
  - Write
  - Delete
  - Modify Permissions
5. Click **OK**.

**Important:** For members of the *administrators* group, the permissions of the newly-created objects will default to Modify Permissions even if they are set to a different level in user preferences.

## Distribution channels

The Distribution Channels option allows the user to specify how they receive notifications. For example, if your company uses IBM SPSS Collaboration and Deployment Services RSS feeds, you may choose to disable email notifications and only enable RSS feeds.

To select distribution channels for notifications and subscriptions:

1. Navigate to the Properties dialog box for the server instance.
2. In the Properties tree, expand the User Preferences list and select **Distribution Channels**.
3. To enable notifications via email, select **Enable Email**. To enable notifications via syndicated feeds, select **Enable RSS**.

Note that the **Enable RSS** option is only available to users who are granted the **Access Syndicated Feeds** action by an administrator.

Syndication feeds provide a format for delivering regularly changing content. IBM SPSS Collaboration and Deployment Services syndicates its notifications as RSS feeds or Atom feeds. An RSS aggregator (feed reader) that supports authenticated feeds is required for viewing notification feeds. For example, RSSBandit and Microsoft® Outlook® 2007 are common desktop-based feed readers, and Google Reader™ and Yahoo® Reader are common Web-based feed readers.

Administrators can override your settings by disabling feeds for all users via the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager. Administrators also define the output format used for all syndication feeds (RSS or Atom), limit the total number of syndication feed entries served to the RSS/Atom subscribers, control which users can access feeds, etc.

4. Click **OK**.

## Subscriptions recipient

The Subscriptions Recipient option allows the user to choose between the following options:

- Typing in an email address to be used as the default for subscriptions
- Using the email address associated with the user from a backing directory, such as active directory

The email address applies per server instance and per user, which means that a user can specify different email addresses for different servers.

To specify an email address for a subscriptions recipient:

1. Navigate to the Properties dialog box for the server instance.
2. In the Properties tree, expand the User Preferences list and select **Subscriptions Recipient**.
3. Select one of the following options:
  - In the **Email Address** field, type the email address to use as the default.
  - Select the **Use email address from directory** check box to use an email address from the directory.
4. Click **OK**.

## Version label (user preferences)

To access server version label properties:

1. In the Content Explorer, right-click on a server name and choose **Properties**. The Server properties dialog box opens.
2. Click + to expand the user preference options.
3. Click **Version Label**. The Version Label dialog box opens.
4. From the **Default label used for files in Jobs** drop-down list, select the default version label to apply to files within a job.
5. Click **OK**.

---

## Version properties

The IBM SPSS Collaboration and Deployment Services Repository allows you to maintain multiple versions of objects.

To access the version listing for an object:

1. In the Properties dialog box, click the Versions tab. A list of versions for the object appears.

The list displays the version number and the created timestamp of the version. The most recent version is labeled *LATEST*. *LATEST* is the version selected by default in the list. The list also displays the *All Versions* entry.

In addition to the general properties defined for IBM SPSS Collaboration and Deployment Services Repository objects, certain properties can also be defined on a version level.

Default version properties include:

- **Description.** A user-defined label for the version.
- **Keywords.** The metadata assigned to IBM SPSS Collaboration and Deployment Services Repository object versions for the purposes of content search.
- **Expiration Date.** The date after which IBM SPSS Collaboration and Deployment Services Repository object versions are no longer active. See the topic “Working with Expiration Dates and Expired Files” on page 31 for more information.

The dialog box also displays custom version properties defined for the server. See the topic “Creating custom properties” on page 39 for more information. The properties are displayed for a selected version in the right pane of the properties dialog box. They can be set for individual versions, for multiple versions, or for all existing versions. The properties allow the user to enhance the metadata for a particular version of an object. For example, the version description property could be used to clearly state its differences from other versions, while specifying keywords makes the search more precise.

To set or modify properties for a single version:

1. Select the version.
2. Modify the properties as necessary. The properties are set to the specified values for the selected version.

To set or modify properties for multiple versions:

1. Hold down the Shift key to select the versions. Any property fields that have different values between the selected versions are empty and marked by alert signs.
2. Modify the properties as necessary. The properties are set to the specified values for the selected versions.

To set or modify properties for all versions:

1. Select *All Versions* in the list. Any property fields that have different values between the existing versions are empty and marked by alert signs.
2. Modify the properties as necessary. The properties are set to the specified values for all versions.

---

## Working with custom properties

Custom properties consist of user-defined metadata that apply to objects within the repository. You must have the appropriate access privileges to create and modify custom properties.

Custom properties are applied per server instance. Custom properties persist across sessions and remain associated with the server instance until they are removed.

Currently, you cannot copy custom properties from one server instance to another. In addition, you cannot make any of the custom properties fields required.

The process for working with custom properties consists of the following tasks:

- Ensuring that you have access privileges to create custom properties
- Creating custom properties
- Setting values for custom properties

Typically, an administrator creates custom properties, and non-administrative users set values for custom properties.

Although custom properties are created at the server level, values for custom properties are set at the content object level. Specifically, the custom properties that you create in the Custom Properties dialog box appear in the General properties dialog box for each content object. See the topic “Editing general properties” on page 26 for more information.

## Verifying access privileges to create custom properties

Before you begin working with custom properties, you must have the appropriate access privileges to create and modify custom properties. If you log in to a server and you do not have access privileges to create custom properties, the custom properties option will not be visible.

A user with administrative privileges to the server instance automatically has permissions to create, edit, and delete custom properties.

Suppose that you log in to a repository and you do not have the appropriate access privileges to create custom properties. If your administrator assigns you access privileges while you are still logged in to the repository, you will need to log out of the repository and then log back into the repository before the updated access privileges take effect.

## Accessing the custom properties dialog box

To access the Custom Properties dialog box:

1. Log in to the server for which you want to create custom properties. See the topic “Logging in to a server” on page 15 for more information.
2. Right-click on the server name and choose **Properties**. The Properties dialog box opens.

*Note:* You must right-click on the server name to view custom properties. If you select any other object within the server folder, only general properties and permissions appear.

3. Click **Custom Properties**. The Custom Properties table appears. If you have previously defined custom properties, they appear in the table.

## Creating custom properties

To create a new custom property:

1. Navigate to the Custom Properties dialog box. See the topic “Accessing the custom properties dialog box” on page 38 for more information.
2. To create a new custom property, click **Add**. The Custom Property Parameters dialog box opens.
3. In the Custom Property Parameters dialog box, you need to provide the following information:
  - **Label.** The label is the name of the custom property as it will appear in the user interface. The label can be up to 128 characters in length and must be unique to the server instance. If you specify a duplicate name, the Duplicate Property dialog box opens indicating that the name is already in use. However, you may specify the same label for custom properties on two different servers.
  - **Property Type.** The property type describes the input value for the custom property. You have the following options:

Table 6. Custom property types.

Property type	Description	Representation in the General properties dialog box
String	Any valid string. String is the default property type.	Text field.
Number	Any numeric value.	Text field.
Single choice	A single selection from the list of options.	Drop-down list.
Multiple select	One or more selections from the list of options.	A series of check boxes.
Yes/No	Either yes or no. Selecting the check box indicates yes. A blank check box connotes no. By default, the check box is cleared. Thus, the default is no.	A single check box.
Date	A valid date in the format m/d/yyyy—for example, 5/18/2006.	A drop-down calendar that prompts the user to specify a date. Alternatively, the user can also type the date.

- **Selection Values.** For custom property types that require the user to choose from a list of values (for example, single select and multiple select), you need to define the values available to the user in the Selection Values list. Items in the list are line delimited. Thus, when you add selection values, each value must appear on a new line.
- **Property Applies to.** This field indicates the objects to which the custom property is applied. You must select at least one check box. You have the following options:

Property Applies to	Description
Folders	Applies to all folders in the repository.
Files	Applies to all files in the repository.
File types	Applies only to the file types you specify. If you select the <b>File types</b> option, the <b>Files</b> option is automatically selected. To specify a file type: <ol style="list-style-type: none"> <li>1. Click the ellipsis button. The File Types dialog box opens.</li> <li>2. From the <b>File Types</b> list, select the file type you want to associate with the custom property. To select multiple file types, press the Ctrl key and select multiple rows. If you select one file type, that file type appears in the text box. If you select multiple file types, the first file type appears in the list followed by an ellipsis.</li> <li>3. Click <b>OK</b>.</li> </ol>
Jobs	Applies only to jobs in the repository.

By default **Folders**, **Files**, and **Jobs** are selected.

- **Property value can be set.** This field indicates whether the property can be set only at the object level or the version level.

4. To save your changes, click **OK**.

## Editing custom properties

After you have created and saved a custom property, you can edit its attributes at any time.

In this context, editing refers to modifying the values available to the user. Editing custom properties takes place at the server level. This process is separate from setting values for the custom property, which takes place at the content object level.

To edit a custom property:

1. Navigate to the Custom Properties dialog box. See the topic “Accessing the custom properties dialog box” on page 38 for more information.
2. In the Custom Properties table, select the custom property that you want to edit.
3. Click **Edit**. The Custom Property Parameters dialog box opens.
4. You can modify any of the parameters in this dialog box except for the following:
  - **Property Type.** When you edit a custom property, the **Property Type** field is disabled. To change the property type, you must create a new custom property. See the topic “Creating custom properties” on page 39 for more information.
  - **Property Value Can Be Set.** The level at which a custom property value is set cannot be modified. When you edit a custom property, this section of the dialog box is disabled.
5. After you have made changes, click **OK** to save them.

## Editing selection values

For the single select and multiple select property types, you can modify the selection values available to users. It is important to note that if you add or delete any of the options in the Selection Values table, that change is propagated to all of the objects on the server associated with that custom property.

For example, suppose you have a multiple select property type with the following options:

- *Harry*
- *Ron*
- *Fred*
- *George*

If you remove *George* from the Selection Values list, the value *George* would be removed from any content objects to which it was previously assigned. In addition, any users that had previously selected *George* would no longer see *George* as an option.

## Searching for custom properties

Custom properties are searchable and used in auditing reports, with the following exceptions:

- **Yes/no label values.** The label that you apply to any yes/no property types will not appear in a search.
- **Numeric values.** For any number property types, the numeric value that you type into the text field will not appear in a search.

See the topic “Searching” on page 17 for more information.

## Deleting custom properties

Custom property definitions are stored in the repository. If you delete a custom property, the custom property and all of its attributes are removed from the repository. Once a custom property is deleted, it cannot be recovered.

It is important to note that when you delete a custom property, the deletion affects all of the objects in the repository that contained that custom property. Once deleted, a custom property will no longer be applied to any object. It is strongly recommended that you do not delete custom properties, unless absolutely necessary.

For example, suppose you have a custom property called **Reviewers** that applies to all of the files on your server. After you delete the custom property and navigate to the Properties dialog box, the **Reviewers** custom property will no longer appear for any file on the server.

To delete a custom property:

1. Navigate to the Custom Properties dialog box. See the topic “Accessing the custom properties dialog box” on page 38 for more information.
2. From the Custom Properties table, select the property that you want to delete. To select multiple custom properties, press the Ctrl key and select multiple rows.
3. Click **Delete**. The Confirm Deletion dialog box opens.
4. Click **OK**. The Custom Properties dialog box reopens, and the Custom Properties table is refreshed.

## Custom Properties and Server Connections

When you delete a server connection, the system removes the connection between the client and the server that hosts your repository. This process does not modify any items in the repository itself. Thus, deleting a server connection has no effect on the custom properties associated with that repository.

---

## Working with topics

Topics allow the definition of a classification system for the content stored in the IBM SPSS Collaboration and Deployment Services Repository, providing a hierarchical map to guide users to the resources they need. Topics function like a directory structure but differ from directories in that a single object can be listed under multiple topics.

For example, you might want to create a topic structure that mirrors your organization, with separate topics for marketing, finance, development, and so on. Users can then choose from the available topics when storing content. In addition, users can limit content searches to specific topics to accelerate the retrieval process. Because a given item can be listed under multiple topics, cross-indexing is also possible.

Alternatively, consider the following topic hierarchy, in which the model type is the basis for classifying resources:

```
Model
  association
    Apriori
    CARMA
    GRI
    Sequence
  clustering
    K-means
    Kohonen
    TwoStep
  decision tree
    C5
    CHAID
    C&RT
    QUEST
  neural network
  screening
    Anomaly detection
```

Feature selection  
statistical  
factor PCA  
linear regression  
logistic regression  
text extraction

Any model in the repository could be assigned a topic in this hierarchy to assist in finding desired resources. For example, a user might want to find all association models that use a specific field. Alternatively, the search may be restricted to CARMA models only.

## Process overview

Topics consist of user-defined metadata that apply to objects within the repository. You must have the appropriate access privileges to create and modify topics. Working with topics consists of a two-part process.

**Creating topic definitions.** Topic definitions and the topic hierarchy are established in the Topic Definitions dialog box. The topics that you create in the Topic Definitions dialog box appear in the General Properties dialog box for each content object. See the topic “Editing general properties” on page 26 for more information.

The following guidelines apply:

- Topic definitions are applied per server instance
- Topic definitions persist across sessions and remain associated with the server instance until they are removed
- Topic definitions cannot be copied from one server instance to another
- Topics that you define cannot be made into required fields

**Assigning topic values.** Although topic definitions are created at the server level, values for topics are set at the content object level. It is important to note that topics are assigned at the object level only. Topics cannot be assigned to folders. In addition, topics cannot be assigned to any items in the *Resource Definitions* folder of the Content Explorer.

Typically, an administrator creates topic definitions, and non-administrative users assign topics.

## Verifying access privileges to create topic definitions

Before you begin working with topic definitions, you must have the appropriate access privileges to create and modify topic definitions. If you log in to a server and you do not have access privileges to create topic definitions, the topic definition option will not be visible.

A user with administrative privileges to the server instance automatically has permissions to create, edit, and delete topic definitions.

Suppose that you log in to a repository and you do not have the appropriate access privileges to create topic definitions. If your administrator assigns you access privileges while you are still logged in to the repository, you will need to log out of the repository and then log back into the repository before the updated access privileges take effect.

## Working with topic definitions

The following tasks can be performed in the Topic Definitions dialog box:

- Create new topic definitions
- Rename topic definitions
- Move topic definitions within the topic hierarchy
- Delete topic definitions



## Accessing the topic definitions dialog box

To access the Topic Definitions dialog box:

1. Log in to the server for which you want to create topics. See the topic “Logging in to a server” on page 15 for more information.

2. Right-click on the server name and choose **Properties**. The Properties dialog box opens.

*Note:* You must right-click on the server name to view topics. If you select any other object within the server folder, only general properties, permissions, and versions appear.

3. Click **Topic Definitions**. The topics hierarchy appears. If you have previously defined topics, they appear in the hierarchy. Otherwise, only the main topics folder appears.

## Creating new topic definitions

To create a new topic definition:

1. Navigate to the Topic Definitions dialog box. See the topic “Accessing the topic definitions dialog box” for more information.

2. Select the folder of the topic hierarchy in which to create the new topic definition.

3. To create a new topic definition, click **Add**. The New Topic dialog box opens.

4. In the **Topic Name** field, type a name for your topic.

5. Click **OK**. The Topic Definitions dialog box reopens and the topic hierarchy is refreshed.

## Renaming topic definitions

Topic definitions can be renamed. After a topic definition is renamed, the change is propagated to all content objects to which the topic is assigned.

Any topic definition can be renamed. However, the root topics directory cannot be renamed. The **Rename** button is disabled when you select the **Topics** folder.

To rename a topic definition:

1. Navigate to the Topic Definitions dialog box. See the topic “Accessing the topic definitions dialog box” for more information.

2. From the topic hierarchy, select the topic you want to rename.

3. Click **Rename**. The Rename Topic dialog box opens.

4. In the **Topic Name** field, type a name for your topic.

5. Click **OK**. The Topic Definitions dialog box reopens, and the topic hierarchy is refreshed.

## Moving topic definitions

Topic definitions can be moved by cutting the topic definition and pasting it to a new location within the topic hierarchy. It is important to note that when a topic is cut from the topic hierarchy, all of topics within it are cut as well.

Any topic definition can be moved. The only exception is the root topics directory. This directory cannot be moved. The **Cut** and **Paste** options are disabled when you select the **Topics** folder.

To move a topic definition:

1. Navigate to the Topic Definitions dialog box. See the topic “Accessing the topic definitions dialog box” for more information.

2. In the topic hierarchy, right-click on the topic you want to move and choose **Cut**.

3. Right-click on the folder into which you want to paste the topic definition and choose **Paste**. The topic definition is moved into its new location.

## Deleting topic definitions

Topic definitions are stored in the repository. If you delete a topic, the topic and all of its attributes are removed from the repository. Once a topic definition is deleted, it cannot be recovered.

For example, suppose you have a topic definition called *Analysis*. After you delete the topic definition and navigate to the Properties dialog box, the *Analysis* topic will no longer appear for any content object on the server.

It is important to note the following about deleting topic definitions:

- When you delete a topic definition, the deletion affects all of the objects in the repository that contained that topic. Once deleted, a topic definition will no longer be applied to any object.
- When you delete a topic definition, all of the topic definitions within it are deleted as well.
- You cannot delete the root topics directory. The **Delete** button is disabled when you select the **Topics** folder.

To delete a topic definition:

1. Navigate to the Topic Definitions dialog box. See the topic “Accessing the topic definitions dialog box” on page 43 for more information.
2. From the topic hierarchy, select the topic that you want to delete.
3. Click **Delete**. The Delete Topic Confirmation dialog box opens.
4. Click **OK**. The Topic Definitions dialog box reopens, and the topic hierarchy is refreshed.

### The effect of deleting server connections on topics:

When you delete a server connection, the system removes the connection between the client and the server that hosts your repository.

This process does not modify any items in the repository itself. Thus, deleting a server connection has no effect on the topic definitions or assignments associated with that repository.

### Deleting topic definitions versus removing topic assignments:

Deleting a topic definition is different from removing a topic assignment.

Deleting a topic definition removes the definition from the system, preventing any content object from being associated with that topic definition. Thus, deleting a topic definition affects all content objects associated with that topic definition.

Removing a topic assignment severs the link between a content object and the topic property. However, removing a topic assignment from a content object does not affect any other content objects that were assigned that topic. See the topic “Removing topics” on page 27 for more information.

## Searching for topics

Topics are searchable and used in auditing reports.

See the topic “Searching” on page 17 for more information.

---

## Bulk property updates

You can update properties for multiple objects simultaneously (in bulk).

The properties that can be updated in bulk include:

- Description
- Keywords
- Topics
- Permissions
- Expiration date
- Custom properties
- Content language

To update properties in bulk:

1. In the Content Explorer, select the objects. Selected objects can include files, folders, and jobs.
2. Right-click and choose **Properties**. The Shared Properties dialog box opens. The dialog box title bar indicates how many objects have been selected.
3. Modify properties on the General, Versions, and Permissions tabs as necessary.

**Note:** The Shared Properties dialog box does not display the Versions tab if any folders are selected for the bulk update.

4. Click **OK**.

## Updating general properties in bulk

To access the general properties for the selected objects:

1. In the Shared Properties dialog box, click **General**.
  - If selected objects are files or jobs, Author, Content Type, Last Modified By, and Modification Date fields are displayed. The dialog box also displays object-level custom properties defined in the repository and the topics common to all objects.
  - If selected objects are folders, only Description, Last Modified By, and Modification Date fields are displayed. The dialog box also displays all custom properties defined for the repository and the topics common to all objects.
  - If selected objects include both files or jobs and folders, only Last Modified By and Modification Date fields are displayed. The dialog box also displays object-level custom properties defined for the repository. The dialog box does not display topics.
  - Any property fields that have different values between the selected objects are empty and marked by alert signs. Green boxes are displayed for check box custom properties when the state of the check box is not common for all selected objects.
2. Modify the property values as necessary. To change the content type of the objects, click the ellipsis button to open the File Types dialog box and select the type. Click **Add** to open the Add Topic dialog box and select topics. To remove a topic, select the topic in the list and click **Remove**. See the topic “Working with topics” on page 41 for more information.

**Note:** If there are jobs among selected objects, the Content Type field is read-only.

## Updating version properties in bulk

To access the version properties for the selected objects:

1. In the Shared Properties dialog box, click **Versions**.

- **Description.** A user-defined label for the version.
- **Keywords.** The metadata assigned to IBM SPSS Collaboration and Deployment Services Repository object versions for the purposes of content search.
- **Expiration Date.** The date after which IBM SPSS Collaboration and Deployment Services Repository object versions are no longer active. See the topic “Working with Expiration Dates and Expired Files” on page 31 for more information.

The dialog box also displays custom version properties defined for the server. See the topic “Creating custom properties” on page 39 for more information. The properties are displayed either for the latest versions of the selected objects or for all versions. Any property fields that have different values between the selected objects are empty and marked by alert signs. Green boxes are displayed for check box custom properties when the state of the check box is not common for all selected objects.

To set or modify properties for the latest versions of selected objects:

1. Select *LATEST* in the list.
2. Modify the properties as necessary.

To set or modify properties for all versions:

1. Select *All Versions* in the list.
2. Modify the properties as necessary.

## Updating permissions in bulk

To access the permissions for the selected objects:

1. In the Shared Properties dialog box, click **Permissions**.

**Owner.** The users who created any of the selected objects are marked with a green box in the Owner column.

**Principal.** The user ID and the number of the selected objects for which the users have permissions defined.

**Permissions.** The Permissions column displays the user's permissions for the selected objects. If the user is not defined in the permissions of all objects, an asterisk (\*) is displayed. To set users permissions to the objects, click in the column and select a permissions level from the list.

**Add.** Add users to the list. See the topic “Adding new users or groups” on page 28 for more information.

**Delete.** Delete users from the list. See the topic “Deleting a user or group from the permissions list” on page 29 for more information.

---

## Chapter 6. Resource definitions

In the Content Explorer, the Resource Definitions folder contains credential definitions, data source definitions, message domains, promotion policies, server definitions, and server cluster specifications. These resources are often required for running jobs .

For example, a IBM SPSS Modeler job step requires the definition of a IBM SPSS Modeler execution server to process the step.

To create new resource definitions, you must have the appropriate actions associated with your login credentials. In addition, you need *Write* permission to the folder that will contain the new definition. Contact your IBM SPSS Collaboration and Deployment Services administrator if your actions or permissions require modifications.

You cannot save any other items into the *Resource Definitions* folder. For example, jobs cannot be saved here. If you try to save or move an item into the *Resource Definitions* folder, an error message appears.

---

### Credentials

Execution of some job steps requires the specification of user credentials, allowing credential validation before executing a job .

User credentials authenticated through an external directory system, such as Active Directory, can be associated with a domain. Write permissions to the Credentials folder are required to create and modify data credentials. See the topic “Modifying permissions for an existing user or group” on page 28 for more information. Also, the user must be assigned the *Define Credentials* action through the role.

### Adding new credentials

To add a new credential:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Credential Definitions* folder.
3. Select the domain in which you want to create a new credential definition.
4. From the File menu, choose:

**New > Credentials Definition**

The Add New Credentials wizard opens.

*Note:* Alternatively, the new credentials dialog box can be accessed by clicking **New** next to the login field on the General tab for certain steps (for example, IBM SPSS Modeler job steps).

5. In the Name field, enter a nickname for the credential. The credential destination dialog box opens.

### Credential destination

If you are defining the credential under a domain, after you have specified the name, the credential destination dialog box opens.

1. Select the domain.
2. Click **Next**. The user and password credentials dialog box opens.

### Specifying a user and password for credentials

After you have specified a name and destination for your credential and clicked **Next**, the User and password credentials dialog box opens.

1. In the user ID field, enter the username for the credential.

2. In the Password field, enter the password that corresponds to the username.
3. In the Confirm password field, reenter the password.
4. If the credential must be validated against a security provider, select the provider using the Security Provider drop-down list. Leave this field blank if provider validation is not necessary for the credential.
5. Click **Finish**. The updated information appears in the *Credentials* folder.

## Creating a new domain

To create a new domain for your credentials:

1. In the Content Explorer, navigate to the *Resource Definitions* folder.
2. Right-click on the *Credential Definitions* folder and select:  
**New > Domain**
3. A new domain appears in the *Credential Definitions* folder, called New Domain.
4. Assign a name to the new domain.

## Server Process Credential

Server Process Credential is the built-in credentials definition of the user profile under which the repository server is run. In Active Directory or OpenLDAP-based single sign-on environment, Server Process Credential can be used instead of regular repository user credentials to:

- Run reporting job steps and schedule time-based jobs
- Query a security provider for a list of user and group profiles

*Important!* Server Process Credential cannot be used in non-single sign-on environments and requires a number of additional configuration steps. For more information, see IBM SPSS Collaboration and Deployment Services Repository configuration documentation.

Server Process Credential cannot be opened, edited, copied, cut pasted, labeled, or versioned, but you can modify permissions to the credential if you have sufficient authority.

---

## Data sources

Open database connectivity (ODBC) provides a mechanism for client programs to access databases or data sources. Similarly, Java database connectivity (JDBC) determines how Java applications access databases.

Data source definitions allow other components of the system, such as IBM SPSS Collaboration and Deployment Services Enterprise View, to access the data sources used in the system. See the topic “IBM SPSS Collaboration and Deployment Services Enterprise View overview” on page 203 for more information.

Data source definitions are created in the Content Explorer and stored in the IBM SPSS Collaboration and Deployment Services Repository. The appropriate permissions are required to create and modify data source definitions. If a user does not have permissions, the **Data Source Definitions** folder is not accessible. See the topic “Modifying permissions” on page 28 for more information. The process for adding a data source definition consists of the following steps:

1. Selecting a data source definition type.
2. Specifying data source definition parameters.

Data source definitions can be exported and imported.

## Selecting a data source definition type

To create a new data source definition:

1. In the Content Explorer, expand the **Resource Definitions** folder. Right-click on the **Data Source Definitions** folder and select:  
**New > Data Source Definition**  
The Select Data Source Definition Type dialog box opens.
2. In the Name field, enter the name for your data source definition. This name, which will appear in the **Data Source Definitions** folder, is used when other components of the system, such as IBM SPSS Collaboration and Deployment Services Enterprise View, invoke the data source definition.
3. From the Type drop-down list, select the data source definition type. Valid values depend on the system configuration and may include *ODBC Data Source*, *JDBC Data Source*, *App Server Data Source*, and *Data Service Data Source*.
4. Click **Next**. The dialog box that appears depends on the data source type that was selected.
  - **ODBC**. The DSN dialog box opens.
  - **JDBC**. The JDBC Name and URL dialog box opens.
  - **App Server Data Source**. The JNDI Name dialog box opens
  - **Data Service Data Source**. The Data Service Data Source Properties dialog box opens.

## Specifying a DSN for an ODBC data source

The data source name (DSN), which corresponds to the database associated with the data source, is specified in the DSN dialog box.

1. In the DSN field, enter the data source name. The name that you enter in this field must match your data source name exactly. The name is also case sensitive.
2. Click **Finish**. The new data source definition appears in the **Data Sources** folder.

## Specifying a JDBC name and URL

The JDBC name and URL, which define how Java applications connect to a database, are specified in the JDBC Name and URL dialog box.

1. In the JDBC Driver Name field, enter the class name of your JDBC driver.
2. In the JDBC Driver URL field, enter JDBC driver URL.
3. Click **Finish**. The new data source definition appears in the **Data Sources** folder.

IBM SPSS Collaboration and Deployment Services is installed with a set of JDBC drivers for all major database systems. It also provides JDBC drivers for IBM SPSS Statistics file access.

For current supported database information, refer to the software product compatibility reports on the IBM Technical Support site.

If single sign-on is enabled for your system, include the optional `AuthenticationMethod` property in the driver URL string to use single sign-on for the database connection. Valid value for this property include *kerberos* and *ntlm* (NT LAN Manager).

## DB2

Class name: *spssoem.jdbc.db2.DB2Driver*

URL string template:

```
jdbc:spssoem:db2://<host>:<port>;DatabaseName=<database>[;AuthenticationMethod=kerberos]
```

## Greenplum

Class name: *org.postgresql.Driver*

URL string template:

`jdbc:postgresql://<host>:<port>/<dbname>`

## Informix

Class name: *spssoem.jdbc.informix.InformixDriver*

URL string template:

`jdbc:spssoem:informix://<host>:<port>;InformixServer=<Informix server>`

## MS SQL Server

Class name: *spssoem.jdbc.sqlserver.SQLServerDriver*

URL string template:

`jdbc:spssoem:sqlserver://<host>:<port>;DatabaseName=<database>[;AuthenticationMethod=kerberos|ntlm]`

## MySQL

Class name: *spssoem.jdbc.mysql.MySQLDriver*

URL string template:

`jdbc:spssoem:mysql://<host>:<port>;DatabaseName=<database>`

## Oracle

Class name: *spssoem.jdbc.oracle.OracleDriver*

URL string template:

`jdbc:spssoem:oracle://<host>:<port>;SID=<database>[;AuthenticationMethod=kerberos]`

## SAP Hana

Class name: *com.sap.db.jdbc.Driver*

URL string template:

`jdbc:sap://<host>:<port>`

## Sybase

Class name: *spssoem.jdbc.sybase.SybaseDriver*

URL string template:

`jdbc:spssoem:sybase://<host>:<port>;databaseName=<database>[;AuthenticationMethod=kerberos]`

## Teradata

Class name: *com.teradata.jdbc.TeraDriver*

URL string template:

`jdbc:teradata://<host>/Database=<dbname>`



## IBM SPSS Statistics data files

Class name: *com.spss.statistics.datafile.jdbc.openaccess.OpenAccessDriver*

URL String template:

```
jdbc:spssstatistics://<hostname>:<port>;ServerDatasource=SAVDB;  
CustomProperties=(<data_file_identifier>;UserMissingIsNull=<1|0>;MissingDoubleValueAsNaN=<1|0>)
```

### Note:

- The *hostname* parameter is the host name or IP address of the computer on which the driver is running. The default value is localhost.
- The default port is 18886.
- Replace *<data\_file\_identifier>* with a string identifying the data file. This may be specified in a variety of ways. Use `CONNECT_STRING=<path_to_file>` to specify the path to the file, where the path is either a repository URI or a path relative to the host on which the driver is running. The path cannot contain any equals sign or semicolon characters. Alternatively, for data files stored in the repository, use `REPOSITORY_ID=<file_id>;REPOSITORY_VERSION=<file_version>` to specify the repository file and version identifiers. The file version may be specified as a label or as a version marker, with spaces replaced by the standard escape character of %20.
- The *UserMissingIsNull* parameter is optional and specifies the treatment of user-defined missing values. 0 indicates that user-defined missing values are read as valid values. 1 indicates user-defined missing values are set to system-missing for numeric variables and blank for string variables. If *UserMissingIsNull* is not specified, it is set to a default value of 1.
- The *MissingDoubleValueAsNaN* parameter is optional and specifies the treatment of missing numeric values. 0 indicates that user missing values are displayed with the original missing value in the data file. 1 indicates that user missing values are read as not a number (NaN). For JDBC, *UserMissingIsNull* should always be set to 1.
- When testing a connection that specifies `AuthenticationMethod=ntlm`, use any credential available in the system. The selected credential is not actually passed to the data source due to the authentication method, but the test requires that a credential be specified.

## System z data sources

IBM SPSS Collaboration and Deployment Services provides the capability to access the following System z data sources using standard DataDirect DB2 drivers:

- DB2 for z/OS
- DB2 LUW on Linux for System z
- Oracle on Linux for System z

## Third-party JDBC drivers

If IBM SPSS Collaboration and Deployment Services does not include a driver for a needed database, you can update your environment to include a third-party driver for the database.

For example, if access to a Netezza or Teradata database is needed, obtain the appropriate driver from the vendor and update your system. To add a JDBC driver to IBM SPSS Collaboration and Deployment Services Deployment Manager:

1. Close IBM SPSS Collaboration and Deployment Services Deployment Manager if it is running.
2. Create a folder named *JDBC* at the root level of the IBM SPSS Collaboration and Deployment Services Deployment Manager installation.
3. Place the driver files in the *JDBC* folder.

Note that for Netezza, the version 5.0 driver should be used to access both version 4.5 and 5.0 databases.

After adding the driver files to your environment, the driver can be used in a data source definition. In the JDBC Name and URL dialog box, type the name and URL for the driver. Consult the vendor documentation for the driver to obtain the correct class name and URL format.

If the data source will be used by the IBM SPSS Collaboration and Deployment Services Repository server, such as in scheduled jobs or reports, the server must also be updated to include the driver files. Contact your administrator to perform the necessary modifications described in the *JDBC drivers* section of the *Installation and Configuration* chapter in the IBM SPSS Collaboration and Deployment Services installation and configuration documentation.

## Specifying a JNDI name for an application server data source

The Java Naming and Directory Interface (JNDI) name identifying the data source is specified in the JNDI dialog box.

1. In the **JNDI Name** field, enter the JNDI name. The name that you enter in this field must match your data source name exactly. The name is also case sensitive.
2. Click **Finish**. The new data source definition appears in the **Data Sources** folder.

## Specifying properties for data service data sources

Data service data source definitions are used to retrieve the data required for scoring tasks. The definitions can specify the structure of context data. Context data are the input data passed with a scoring request in real time.

For example, when a score is requested for a customer based on credit rating and geocode, the credit score and geocode will be the context data for the request. Data service data source definitions can also be used retrieve data from nonstandard sources using custom drivers.

Properties for data service data sources are specified in the Data Service Data Source Properties dialog.

To specify data service data source properties:

1. Specify whether to use context data or a custom driver.
  - Context data.** If the context data option is selected, tables can also be defined.
  - Custom driver.** If the custom driver option is selected, a driver class must be provided. Additional driver properties can also be added. For information on creating custom drivers, see IBM SPSS Collaboration and Deployment Services customization documentation.
2. Optionally, tables can be defined for both the context data and custom driver options.
3. Click **Finish**. The new data source definition appears in the **Data Sources** folder.

## Defining tables for data service data sources

Optionally, tables can be defined or imported for both the context data and custom driver options.

To define tables:

1. After the context data or custom driver parameters have been specified, click **Next**. The Define Tables dialog box appears.
  - Add.** Click to define a table. See the topic "Table properties" on page 53 for more information.
  - Edit.** Click to edit the selected tables. See the topic "Table properties" on page 53 for more information.
  - Import.** Click to import tables from previously exported data structure definition.
  - Export.** Click to export data structure definition.
2. After data structure information has been provided, click **OK**.
3. Click **Finish**. The new data source definition appears in the **Data Sources** folder.

## Table properties

To define a table for a data service data source:

1. In the Define Tables dialog, click **Add** or select a table definition and click **Edit**. The Table Properties dialog appears.
  - Name.** Define the table name.
  - Add Column.** Click to add a column. The Add Column dialog appears. See the topic “Adding columns” for more information.
  - Manage Keys.** Click to define keys for the table. See the topic “Managing keys” for more information.
  - Remove Column.** Click to remove a selected column.
2. After the table information has been provided, click **OK**. The table appears on in the list in the Define Tables dialog.

### Adding columns:

To add a column to a table:

1. In the Table Properties dialog, click **Add Column**. The Add Column dialog appears.
  - Name.** Define the column name.
  - Type.** Select the column data type from the drop-down list, for example, Boolean, Date, Decimal, String, or Timestamp.
2. After the column information has been provided, click **OK**. The column appears on in the list in the Define Tables dialog.

### Managing keys:

The Manage Keys dialog box provides options for viewing and defining keys for each table. A **key** is a set of columns that can be used to identify or access a particular row or rows. The key is identified in the description of a table, index, or referential constraint. The same column can be part of more than one key.

**Note:** A key can be based on more than one column. Rows in two tables match if the values of all columns in the table keys match each other, in the same order.

A unique key is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain NULL values. For example, an employee number column can be defined as a unique key, because each value in the column identifies only one employee. No two employees can have the same employee number.

### Assigning keys to a table

1. Click **Manage Keys**. The Manage Keys dialog box opens.
2. Click **Add** to define a new key. The Key Properties dialog box displays, allowing you to define or edit an existing key.
3. Enter an appropriate key name in the **Name** field and click **Unique** if the key is unique. This name appears under the list of keys and is the name a consuming application will use to determine the key.
4. Select the appropriate columns from the Available Columns table and click the right-arrow button. The selected column names are displayed in the Key Columns list. Derived columns can be used as keys.

You can remove existing columns from the Key Columns list by selecting the appropriate columns and clicking the left-arrow button.
5. Click the **Up** and **Down** buttons to rearrange the key column order.
6. Click **OK** to save the key definition. The Key Properties dialog box closes, and the key name is displayed in the Manage Keys dialog.

7. Click **Close**. The selected table displays the number of currently defined keys in the corresponding column.

## Modifying data source definitions

After a data source definition has been established, definition properties can be modified. For example, a JDBC driver URL may need to be updated.

Although data source parameters can be modified, the data source definition type cannot be changed. An ODBC data source cannot be converted into a JDBC data source, and vice versa. To change the data source definition type, a new data source definition must be created.

To modify the properties of an existing data source definition:

1. Navigate to the **Data Source Definitions** folder.
2. Double-click on the data source definition to be modified. The dialog box that opens depends on the data source definition type.
  - **ODBC**. The DSN dialog box opens.
  - **JDBC**. The JDBC Name and URL dialog box opens.
3. Make the modifications.
4. Click **Finish**. The changes are applied to the data source definition.

---

## Message domains

Message domains define JMS (Java Messaging Service) queues or topics used by IBM SPSS Collaboration and Deployment Services to communicate with third-party applications.

Messages domains are used to set up message-based schedules and message-based job steps for IBM SPSS Collaboration and Deployment Services jobs. Message domain definitions reference the messaging service which is set up outside IBM SPSS Collaboration and Deployment Services, usually using the JMS infrastructure of the underlying application server.

Write permissions to the Message Domains folder are required to create and modify data credentials. See the topic “Modifying permissions for an existing user or group” on page 28 for more information. Also, the user must be assigned the *Define Message Domains* action through the role.

IBM SPSS Collaboration and Deployment Services Deployment Manager provides the ability to create new message domains and modify and delete domain definitions.

## Creating a new message domain

To add a message domain definition:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Message Domains* folder.
3. From the File menu, choose:

**New > Message Domain**

The Add New Message Domain wizard opens.

The wizard provides the capability to specify new message domain name and JMS properties.

### Message domain name

The first step of adding a new message domain involves the definition of the domain name.

1. In the **Name** field, enter the name for the message domain.

*Note:* Currently the only supported message domain type is topic.

2. Click **Next**. The Message Domain Page of the wizard opens.

## Message domain properties

1. On Message Domain Properties dialog, specify the following JMS properties:
  - **Destination Name.** The name of the topic or queue.
  - **Credentials.** Optional credentials. The credentials will only be required in certain instances, depending upon the configuration of the JMS server and how we are required to connect to the JMS server (whether the JMS Message topic is secured or not).

*Note:* If JBoss JMS messaging service is being used, message domain must be specified in the topic/<topic name> format.

2. Click **Finish**. The updated information appears in the Message Domains folder.

## Modifying message domain definitions

To modify a message domain definition:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Message Domains* folder.
3. Double-click the message domain. The Edit Message Domain dialog opens.
4. Modify message domain properties as necessary.
  - **Destination Name.** The name of the topic or queue.
  - **Credentials.** Optional credentials. The credentials will only be required in certain instances, depending upon the configuration of the JMS server and how we are required to connect to the JMS server (whether the JMS Message topic is secured or not).
5. Click **Finish**. The updated information appears in the Message Domains folder.

---

## Promotion policies

A promotion policy is a resource definition that specifies rules and properties to apply when promoting objects. Use of a promotion policy prevents having to redefine the set of rules each time an object is promoted. Instead, you apply the policy to the promotion request, and the policy rules are automatically enforced. A promotion policy is defined once, but can be applied to any number of promotion requests. Information specified in a promotion policy includes the following:

- promotion timing
- resource definition inclusion
- related resource exclusion

To create or modify a promotion policy, the role for the user must include the *Define Promotion Policies* action. In addition, the user must have write permissions to the *Promotion Policies* subfolder of the *Resource Definitions* folder.

IBM SPSS Collaboration and Deployment Services Deployment Manager provides the ability to add promotion, modify, and delete promotion policy definitions.

## Adding promotion policies

To add a promotion policy:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Promotion Policies* folder.
3. From the File menu, choose:

**New > Promotion Policy**

The Add Promotion Policy wizard opens.

## Promotion timing

Promotion timing determines when the promoted object becomes available in a new IBM SPSS Collaboration and Deployment Services Repository. In immediate promotion, the promoted object is added to the new repository server automatically as part of the promotion request. In this case, the promotion policy must specify the target repository server as well as valid credentials for accessing that server to enable the source and target repository servers to communicate with each other. In contrast, delayed promotion saves the promoted object to a specified file on the file system for subsequent transfer at a later time. In this case, the promotion policy includes no information about the target repository server. The promotion is completed by manually importing the resulting file into the target server.

The Promotion Timing page of the wizard allows you to specify whether the promotion should be immediate or delayed. The following options are available:

- **Immediate.** Select the option to immediately promote the objects to a different repository server.
- **Target Server.** If you selected the Immediate option, select the target repository server.
- **Target Credentials.** If you selected the Immediate option, select the credentials for accessing the target repository server. For information about credentials, see “Adding new credentials” on page 47. In addition, the creation of new items in the target server may require that the role associated with the target credentials include certain actions. See the topic “Promotion considerations” on page 69 for more information.
- **Delayed.** Select the option to save to an export file which can be later imported into the targeted server later.

1. Click **Next**. The Resource Definition Policy page of the wizard opens.

## Promotion policy resource definitions

The resource definition portion of a promotion policy specifies how to handle the resource definitions referenced by a promoted object. For example, a job may include steps that refer to server and credential definitions. Some steps may rely on data source definitions or on a version of the Enterprise View. If these definitions do not exist in the target server, the promoted item will not function properly without manual intervention to redefine these properties. By including the necessary resource definitions in the promotion, the promoted item should function in the target server exactly as it does in the source server without any redefinition. However, if you are sure the necessary resource definitions exist in the target server, the definitions from the source server can be omitted from the promotion.

The Resource Definition Policy page of the wizard allow you to specify the rules for resource definitions handling.

1. Select one of the following options:

- **Recommended** . Promote resource definitions where no identifier or name conflicts exist in the target repository.
- **Exclude.** Do not promote resource definitions.
- **Include.** Promote resource definitions.

2. If Include option is selected, you can specify which of the following resource types should be excluded from promotion:

- Credentials
- Data sources
- Message domains
- Servers
- Topics
- Custom properties

- Enterprise View
- Notifications
- Server clusters
- Promotion policies

3. Click **Finish** to add the promotion policy or click **Next** to proceed to MIME Type Filter page.

## MIME type filtering

If an object being promoted uses other IBM SPSS Collaboration and Deployment Services Repository items, the promotion policy should include information about which of those items should be excluded from the promotion. For example, when promoting IBM Analytical Decision Management objects, it may be necessary to include a version of an associated rule file for a IBM SPSS Modeler stream but exclude the Data Provider Definition version used. The promotion policy specifies the items to exclude by their MIME types. In the previous example, the policy would specify *application/x-vnd.spss-data-provider* as a MIME type to exclude from the promotion.

The MIME Type Filter page of the wizard allow you to optionally specify the content types of resources used by the promoted item that should be excluded from the promotion.

1. Select the MIME type of objects to be excluded from promotion:
  - To add a file type, click **add**. The File Types dialog box opens. Click the file type entries to select or clear. Use the Ctrl or Shift keys to select multiple entries. After all file types have been selected, click **OK**.
  - To delete a file type, select the file type and click **Delete**.
2. Click **Finish**.

## Modifying promotion policies

To modify a promotion policy:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Promotion Policies* folder.
3. Right-click a promotion policy and select **Open**.

The Edit Promotion Policy wizard opens. Specify the policy properties such as promotion timing, resource definition handling, and MIME types of used resources to be excluded from promotion.

## Deleting promotion policies

To delete a promotion policy:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Promotion Policies* folder.
3. Right-click a promotion policy and select **Remove**.

---

## Server definitions

Executing an IBM SPSS Collaboration and Deployment Services Repository resource as a job step requires the specification of an appropriate corresponding server to process the instructions contained in the job step. The connection information for such a server is specified within a server definition.

Server definitions can be classified as either execution servers or repository servers.

- Execution servers process the contents of an IBM SPSS Collaboration and Deployment Services Repository resource. The execution server type must correspond to the resource type being processed. A SAS job step, for example, requires a SAS server definition.

- A repository server corresponds to an IBM SPSS Collaboration and Deployment Services repository installation. A server of this type is typically used by job steps that need to return result artifacts to a repository.

Server definitions are contained in the *Resource Definitions* folder of the Content Explorer. Specifically, they are defined in the *Servers* subfolder.

## Adding new server definitions

To add a new server:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Click the *Servers* folder.
3. From the File menu, choose:

### New > Server Definition

The Add New Server Definition wizard opens. Alternatively, the new server definition dialog box can be accessed by clicking **New** next to a server field on the General tab for some steps. The process for defining new servers consists of:

1. Naming the server definition and specifying its type. Note that the server types available depend on which product adapters are installed to the repository.
2. Selecting a location in the *Servers* folder for the definition.
3. Specifying parameters for the server that define connection or execution information. The parameter set depends on the server type.

### Server definition type

The Server Definition Type page specifies identification information for a server definition.

1. In the Name field, enter the name that you want to assign to your server definition.
2. From the Type drop-down list, select the server type. Note that the server types available depend on which product adapters are installed to the repository.

*Note:* If the server is to be used to run IBM SPSS Statistics custom dialogs in IBM SPSS Collaboration and Deployment Services Deployment Portal, then select **Remote IBM SPSS Statistics Server** for the server type.

Click **Next** to define additional parameters.

### Server destination

The Server Destination page specifies the location in the IBM SPSS Collaboration and Deployment Services Repository in which the definition is stored.

1. Navigate to the folder, and select it.
2. Click **Next**. The next dialog box that opens depends on the server type that you selected.

### Content repository server definition

A Content Repository Server Definition specifies the connection parameters for a IBM SPSS Collaboration and Deployment Services Repository server.

1. In the **Server URL** field, enter the complete connection URL for the server.

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

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

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

## 2. Click **Finish**.

The new definition is included in the Servers folder.

### SAS server parameters

A SAS Server Definition specifies the execution file used for processing SAS job steps.

The execution file must be installed on the same host as the IBM SPSS Collaboration and Deployment Services Repository. Alternatively, a remote process server could be installed on a remote machine to allow IBM SPSS Collaboration and Deployment Services to access an execution file installed on the remote machine.

In the Executable field, specify the full path to the *sas.exe* file to use as the execution server. If the system path includes the location of this file, the full path can be omitted and the default value of *sas* used.

On some systems, spaces in the path for the execution file may cause failures for jobs using the execution server. These problems can be eliminated by specifying the path using 8.3 notation. For example, instead of:

C:\Program Files\SAS Institute\SAS\V8\sas.exe

Specify:

C:\Progra~1\SASINS~1\sas\v8\sas.exe

After defining the executable path, click **Finish**. The new definition appears in the *Servers* folder.

### Remote process server parameters

A Remote Process Server Definition specifies the connection parameters for a server configured for remote processing. A remote process server allows IBM SPSS Collaboration and Deployment Services to access functionality installed on the remote machine.

1. In the Host field, enter the name of the host where the remote process server resides.
2. In the Port field, enter the port number to be used to connect to the host.
3. To specify that Secure Socket Layer (SSL) is to be used for the server connection, select **This is a secure port**.
4. Click **Finish**. The new definition appears in the *Servers* folder.

### Modifying server definitions

To modify a server definition:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Open the *Servers* folder.
3. Double-click the server to be modified. The Edit Server Definition dialog opens.

4. Modify the server definition parameters as necessary.
5. Click Finish to save your changes.

---

## Server clusters

IBM SPSS Statistics, IBM SPSS Modeler, and remote process servers can be grouped into server clusters to allow load balancing across servers.

When a job step uses a server cluster for execution, IBM SPSS Collaboration and Deployment Services determines which managed server in the cluster is best suited to handle processing requests at that time and routes the request to that server.

The load balancing algorithm for routing server requests uses a weighted least-connection algorithm based on server scores and server loads. When a new connection to a cluster is requested, the system determines a score for each running server in the cluster by using the following formula:

$$W_i * C_i / (N_i + 1)$$

The value of  $W_i$  is the weight associated with server  $i$ . The  $C_i$  value is the number of CPUs for server  $i$ . The value of  $N_i$  is the number of current and pending connections for server  $i$ .

Using the average server load, the system classifies each server as either *available* or *busy*. The new connection is assigned to the *available* server having the highest score. If there are no *available* servers, the connection is assigned to the *busy* server having the highest score. If multiple servers have the same score, the connection is assigned to the one that has the smallest server load.

If two servers in a cluster have the same number of CPUs and server load, the ratio of the number of connections for each will depend entirely on the server weights. For example, if server A has a weight that is twice the weight of server B, server A will handle twice the number of connections. Conversely, if two servers have the same weight and server load, the ratio of the number of connections for each will depend entirely on the number of CPUs for the servers. If server C has eight CPUs and server D has two, server C will handle four times the number of connections.

Notice that the server scores are based on both the number of current and pending connections. If many job steps initiate connections to a server cluster simultaneously, a server in the cluster may not be able to report new connections as current before another connection request is attempted. By including the pending connection count, the scores accurately reflect the impending server load allowing the algorithm to optimize the distribution of requests across all servers in the cluster. A configuration setting defines the time interval during which a connection is classified as pending. For information on modifying this value, consult the administrator documentation.

Server cluster definitions are created in the Content Explorer and stored in the IBM SPSS Collaboration and Deployment Services Repository. Write permissions to the Server Clusters folder are required to create and modify server clusters. See the topic "Modifying permissions for an existing user or group" on page 28 for more information. Also, the user's role must include the *Define Server Clusters* action.

IBM SPSS Collaboration and Deployment Services Deployment Manager provides the ability to create new server clusters and modify and delete cluster definitions.

## Creating a new server cluster

To add a server cluster definition:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Select the *Server Clusters* folder.
3. From the File menu, choose:

## New > Server Cluster Definition

The Add New Server Cluster wizard opens. This wizard provides the capability to specify the cluster name and settings.

### Server cluster definition name

The first step of adding a new server cluster involves the definition of the cluster name.

1. In the **Name** field, enter the name for the cluster.
2. Click **Next**. The Server Cluster Settings page of the wizard opens.

### Server cluster settings

The settings for a server cluster define the servers included in the cluster and the weight associated with each server.

To add a server to the cluster, click the Add button. The Add Servers to Cluster dialog box appears. See the topic “Add servers to cluster” for more information.

To remove a server from the cluster, select the server to be removed in the server list and click the Remove button. To select multiple servers for removal, press the Ctrl key when selecting servers.

To modify the weight for a server, click the weight to be modified and click the ellipsis button that appears in the cell. The Set Server Weight dialog box appears. See the topic “Set server weight” for more information.

After defining the cluster settings, click **Finish**.

#### Add servers to cluster:

The Add Servers to Cluster dialog box lists all servers on the network currently registered with IBM SPSS Collaboration and Deployment Services.

Select the server to be added to the cluster and click the OK button. To select multiple servers, press the Ctrl key when selecting servers.

#### Set server weight:

The Set Server Weight dialog box allows the specification of a weight for the server.

Weights must be between 1 and 100, inclusive. Weights represent a relative value for workload capacity. For example, a server with a weight of 10 has 10 times the capacity of a server with a weight of 1. Click OK to apply the weight value to the server and return to the Server Cluster settings page.

## Modifying a server cluster

To edit a server cluster definition:

1. In the Content Explorer, open the *Resource Definitions* folder.
2. Select the *Server Clusters* folder.
3. Right-click the cluster to be modified and select Open.

The Edit Server Cluster dialog box opens. This dialog box provides the capability to modify the cluster settings. See the topic “Server cluster settings” for more information.

---

## Importing resource definitions

Like other IBM SPSS Collaboration and Deployment Services Repository objects, resource definitions can be exported and imported.

The procedure for exporting resource definitions is similar to the procedure for exporting regular folders. See the topic “Exporting folders” on page 64 for more information. Note that all existing credential, data source, and server information is included in the export file.

Resource definitions in a IBM SPSS Collaboration and Deployment Services export file can be imported "wholesale" or separately. The procedure for importing all resource definitions from an export file is similar to the procedure for regular folder import. See the topic “Importing folders” on page 65 for more information.

To import resource definitions separately:

1. In the Content Explorer, expand the *Resource Definitions* folder.
2. Right-click on a resource definition and select **Import**. The Import Folder dialog box opens.
3. Navigate to the .pes import file and select it.
4. Click **Open**. The progress dialog box appears. When the import is complete, the Import dialog box opens.
5. Click **OK**. The imported folder and its contents appear in the Content Explorer tree.

**Note:** The import process performs a check of execution server definitions. If one of the definitions is invalid, the entire process fails.

---

## Chapter 7. Export, import, and promotion

---

### Overview

If you maintain multiple IBM SPSS Collaboration and Deployment Services Repository instances, you may have a need to move items from one repository server to another.

For example, you may have a server dedicated to objects under current development. When an object is ready, you add it to a test server to evaluate its performance. If the object performance meets the testing criteria, you add the item to a production server for use within your enterprise. IBM SPSS Collaboration and Deployment Services provides the following ways to transfer repository objects:

- **Export** can be used to save the content of an entire repository folder to a compressed archive file with the extension .pes. The archive will include all subfolders, child objects, and any custom properties applied to the folder. It is also possible to include external references, that is, objects outside the exported folder, such as resources definitions, in the export. Objects can be selected for export based on label. When the export file is **imported** into the target repository, the directory structure of the source folder will be re-created at the root of the folder into which it is imported. See the topic “Exporting folders” on page 64 for more information.
- **Promotion** allows you to transfer individual repository objects with their dependent resources using **promotion policies**. See the topic “Promotion” on page 69 for more information.

**Note:** When you use export, import, or promotion to transfer objects between repository servers, the servers must all be the same version and have the same adapters installed. If you need to move objects between repository servers that are of different versions, follow the migration process described in the repository installation documentation.

### Job components that are migrated during export and import

When importing and exporting jobs, the following components of a job are migrated:

- Versions
- Notifications
- Job schedules

The job history associated with a job is not migrated during the import or export process.

### Exporting external references

External references are resources that are defined outside of the folder that is being exported. For example, external references may consist of IBM SPSS Modeler streams and other resources (e.g., syntax files) that defined in different folder (i.e., not in the folder selected for export).

By default, external references are included during the export process. However, the **Include external references** option can be disabled when exporting objects. See the topic “Exporting folders” on page 64 for more information.

It is important to note that selecting the **Include external references** check box has no effect in the following instances because external references are not applicable.

- When exporting from the root of the Content Repository tree
- When exporting the Resource Definitions folder

### Export and import restrictions

When exporting and importing, the system enforces certain restrictions.

## Export Restrictions

The content of the *Submitted Jobs* folder cannot be exported.

## Import Restrictions

- All jobs containing streams must have the streams in the same relative location on the target server unless the streams are included in the .pes file.
- The environment of the target IBM SPSS Collaboration and Deployment Services Repository must include any adapters needed for working with the artifacts in the archive being imported. For example, if the archive contains IBM SPSS Modeler server definitions, the target environment must include the IBM SPSS Modeler adapters.
- When importing large archive files, it may be necessary to use a disk mapping implementation for identifier lookup. Use the IBM SPSS Collaboration and Deployment Services Deployment Manager to change the *Resource Transfer Lookup Table* configuration value to *DISK* before attempting the import. For more information, see the Administrator's Guide.

## Recommended import order

When importing content into IBM SPSS Collaboration and Deployment Services Repository, it is recommended that the following order be followed:

- For IBM SPSS Collaboration and Deployment Services Enterprise View users, import Enterprise View objects (usually performed once). See the topic “Importing and exporting IBM SPSS Collaboration and Deployment Services Enterprise View objects” on page 253 for more information.
- Import IBM SPSS Collaboration and Deployment Services Resource Definitions (usually performed once). See the topic “Importing resource definitions” on page 62 for more information.
- Import other objects that will be stored in the IBM SPSS Collaboration and Deployment Services Repository (usually performed multiple times).

## Security permissions when importing folders

The following security restrictions apply when importing folders:

- The system accepts the security configuration (permissions) of the parent folder. See the topic “Modifying permissions” on page 28 for more information.
- If a job contains a server-related job step, the system tries to match the server and credential names between the source and target machines. If the server and credential definitions do not match, the import process will fail. See the topic “Adding new credentials” on page 47 for more information.

---

## Exporting folders

The process for exporting files consists of the following tasks:

1. Specifying a location for the exported file.
2. Determining whether to export external references.

To export a folder:

1. In the Content Explorer, right-click on the folder to export, and select **Export**. The Export dialog box appears.
2. To select an export destination, click **Browse**. The Browse for Folder dialog box opens.
3. Select the location in which you want to store the exported folder. An existing folder can be selected or a new folder can be created.
  - **Using an existing folder.** Navigate to an existing folder within the tree, and select it.
  - **Creating a new folder.** To create a new folder, click **Make New Folder**. A new folder is inserted into the tree with the name *New Folder*. Rename this folder.

4. Select one of the following options for including versions of child in the export.
    - **Export all labeled versions.** Include all labeled versions.
    - **Export all versions** Include all versions (labeled and unlabeled).
    - **Export labeled versions** Include versions with the specific labels.
  5. If export labeled versions option is selected, specify the labels to be included in export.
    - To add a label to the list of export labels, select an available label and click **Add**.
    - To add all labels to the list of export labels, click **Add All**.
    - To remove a label from the list of export labels, select a label and click **Remove Label**.
    - To remove all labels from the list of export labels, click **Remove All**.
  6. Click **Save**. The Export dialog reappears.
  7. Determine whether to include external references in the exported file. By default, external references are included during the export process. To exclude external references, clear the **Include external references** check box. See the topic “Exporting external references” on page 63 for more information.
  8. Click **OK**. The progress dialog box appears. When the export is complete, the Export dialog box opens.
  9. Click **OK**.
- 

## Importing folders

The process for importing files consists of the following tasks:

1. Specifying a path.
2. Determining how the system should resolve import conflicts.

*Important:* The version import is processed differently for IBM SPSS Collaboration and Deployment Services Enterprise View objects. See the topic “Importing and exporting IBM SPSS Collaboration and Deployment Services Enterprise View objects” on page 253 for more information.

To import a file:

1. In the Content Explorer, right-click on the directory into which you want to import the folder and select **Import**. The Import dialog box opens.
2. Click **Browse** to navigate to the file that you want to import . The Browse for Folder dialog appears.
3. Select the file to import, and click **Open**. The Import dialog returns. The name of the files selected for import appears in the Import File field. The version product version number associated with the imported file appears in the Import file version field. The import file version is determined by the system and is not modifiable.
4. Specify whether the system should resolve import conflicts individually or globally by selecting the applicable option. By default, IBM SPSS Collaboration and Deployment Services Deployment Manager resolves conflicts individually. If global conflict resolution is selected more options are enabled in the Import dialog. Specify the settings for global conflict resolution. See the topic “Resolving import conflicts globally” on page 66 for more information.
5. Click **OK** to launch the import process.
6. If any conflicts arise during the import process, the Import Conflicts dialog appears. See the topic “Resolving import conflicts” on page 66 for more information.
7. If no conflicts occurred, the progress dialog box appears. When the import is complete, the Import complete dialog box opens.
8. Click **OK**. The imported folder and its contents appear in the Content Explorer tree.

## Resolving import conflicts

At times, an object that is imported may conflict with an existing object in the repository. For example, the imported object may be a duplicate of an object that is currently in the repository. When conflicts occur, they must be resolved before the import process can continue.

Conflicts can be resolved in a global or individual basis. Global conflict resolution is done by the system. Individual conflict resolution requires a user-specified decision for each conflict.

The method used to resolve conflicts is specified in the Import dialog. By default, the system resolves conflicts individually. Any setting changes made in this dialog persist until they are modified in a future session.

### Resolving import conflicts globally

If Resolve conflicts globally is selected in the **Import** dialog box, additional settings must be specified. These settings determine how the system resolves conflicts that arise during the import process. Unlike individual conflict resolution, a list of conflicts prompting user decisions for each conflict will not appear if global conflict resolution is selected. The system will automatically deal with the conflicts that arise based on the criteria specified.

The global resolution options are organized as follows:

**Resolve Resource Conflicts.** If duplicate IDs or names are encountered during the import process, the following options are available:

- *Keep target item.* The target item will be maintained. The source item with the duplicate ID, which is contained in the .pes file, will be ignored.
- *Add new version of target item or rename source item.* This option is typically used to resolve ID conflicts or naming conflicts. If a duplicate ID conflict occurs between the source object and the target object, then a new version of the object will be created in the target location. If a naming conflict occurs, the imported object will be renamed in the target location. Typically, renamed objects are appended with *\_1*, *\_2*, and so forth. In the event that two versions of an object have the same label, the system keeps one label and discards the duplicate label because no two versions of the same item can have the same label. Whether the source label or the target label is maintained depends upon the value specified in the *Use labels from* drop-down field. By default, the label in the source object will be maintained when a conflict occurs. To use labels from the target directory, select **Target** from the *Use labels from* drop-down list.

**Lock Resolution.** Locked resources can affect the import process. See the topic “Object Locking” on page 21 for more information. When setting import defaults for lock resolution, the following options are available:

- *Continue import, even if some objects cannot be imported due to locking conflicts.* Selecting this option may result in a partial import.
- *Abort import if all objects cannot be imported.* If any conflicts are encountered due to object locks, the import process will be terminated and fail.

**Resolve Invalid Version Conflicts.** If an invalid version is encountered during the import process, the following options are available:

- *Import.* The invalid version will be imported. This option is the default.
- *Discard.* The invalid version will be deleted.

**Resource Definitions.** The system imports resource definitions by using one of the following rules:

- **Recommended.** A resource definition is imported only if the identifier or name does not conflict with a target definition. Any resource definitions that have a conflict are not imported.
- **Exclude.** No resource definitions from the import files are imported. Imported objects may need to be modified to reference available resource definitions.



- **Include.** All resource definitions in the import file are imported. You can select one or more resource definition types to exclude from the import by selecting the corresponding check box.

## Troubleshooting Import Failures Caused By Duplicate Name Conflicts

Occasionally, the import process may fail due to import configuration settings. If the import failure was caused by naming conflicts, it is possible that changing a setting in the Import dialog may resolve the issue. To resolve an import failure due to naming conflicts:

1. Repeat the import process. See the topic “Importing folders” on page 65 for more information.
2. In the *Resolve Resource Conflicts* section of the Import dialog, select the **Add new version of target item or rename source item** option.

## Resolving import conflicts individually

If *Resolve conflicts individually* was selected in the *Import* dialog and conflicts arise during the process, the Import Conflicts dialog appears.

The process for resolving individual conflicts consists of the following tasks:

1. Specifying conflict resolution defaults.
2. Determining if defaults need to be overridden in the import conflicts table.

### Individual conflict resolution defaults:

Defaults for the following import conflicts need to be specified:

**Resolve Resource Conflicts.** If duplicate IDs or names are encountered during the import process, the following options are available:

- *Keep target item.* The target item will be maintained. The source item with the duplicate ID, which is contained in the *.pes* file, will be ignored.
- *Add new version of target item or rename source item.* This option is typically used to resolve ID conflicts or naming conflicts. If a duplicate ID conflict occurs between the source object and the target object, then a new version of the object will be created in the target location. If a naming conflict occurs, the imported object will be renamed in the target location. Typically, renamed objects are appended with *\_1*, *\_2*, and so forth. In the event that two versions of an object have the same label, the system keeps one label and discards the duplicate label because no two versions of the same item can have the same label. Whether the source label or the target label is maintained depends upon the value specified in the *Use labels from* drop-down field. By default, the label in the source object will be maintained when a conflict occurs. To use labels from the target directory, select **Target** from the *Use labels from* drop-down list.

**Lock Resolution.** Locked resources can affect the import process. See the topic “Object Locking” on page 21 for more information. When setting import defaults for lock resolution, the following options are available:

- *Continue import, even if some objects cannot be imported due to locking conflicts.* Selecting this option may result in a partial import.
- *Abort import if all objects cannot be imported.* If any conflicts are encountered due to object locks, the import process will be terminated and fail.

**Resolve Invalid Version Conflicts.** If an invalid version is encountered during the import process, the following options are available:

- *Import.* The invalid version will be imported. This option is the default.
- *Discard.* The invalid version will be deleted.

**Resource Definitions.** The system imports resource definitions by using one of the following rules:

- **Recommended.** A resource definition is imported only if the identifier or name does not conflict with a target definition. Any resource definitions that have a conflict are not imported.
- **Exclude.** No resource definitions from the import files are imported. Imported objects may need to be modified to reference available resource definitions.
- **Include.** All resource definitions in the import file are imported. You can select one or more resource definition types to exclude from the import by selecting the corresponding check box.

### Troubleshooting Import Failures Caused By Duplicate Name Conflicts

Occasionally, the import process may fail due to import configuration settings. If the import failure was caused by naming conflicts, it is possible that changing a setting in the Import dialog may resolve the issue. To resolve an import failure due to naming conflicts:

1. Repeat the import process. See the topic “Importing folders” on page 65 for more information.
2. In the *Resolve Resource Conflicts* section of the Import dialog, select the **Add new version of target item or rename source item** option.

#### Working with the import conflicts table:

Individual import conflicts are listed in the Import Conflicts table.

The Import Conflicts table contains the following information. With the exception of the Action column, none of the information in the table can be modified.

**Source Path.** The source location of the object to be imported.

**Hierarchy.** The location within the repository hierarchy--for example, folder.

**Conflict.** The nature of the import conflict. Examples include *Duplicate* and *Invalid version*.

**Target.** The destination location of the object to be imported. The Destination path field is populated only for duplicates.

**Marker.** The object identifier of the item. Conflicts are determined by this identifier. For example, the system searches for duplicate object identifiers.

**Additional information.** Any supplementary information about the import conflict. For example, *Missing IBM SPSS Modeler stream* might appear in this column, if that is the cause of an invalid version.

**Action.** The *Default* that appears in the Action column refers to the default actions specified in the *Conflict Resolution Defaults* section of this dialog box. In Action column, the default can be overridden on an individual conflict basis. The options that appear vary depending upon the type of conflict. For example:

- Actions for duplicates include *Use global default*, *Keep target item*, and *Use imported item*.
- Actions for invalid versions include *Default*, *Import*, and *Ignore*.

To modify the default Action setting:

1. Highlight the conflict whose default setting you want to modify.
2. Click the ellipsis button next to *Default* in the Action column. A dialog, which is specific to the type of conflict appears.
3. Select the options within the dialog box and click **OK**. The import conflicts table returns. For items whose default actions were modified, *Custom* appears in the Action column.
4. After all default actions have been modified in the table, click **OK** to continue the import process.

---

## Promotion

Promotion provides a policy-based approach to transferring individual objects between IBM SPSS Collaboration and Deployment Services Repository instances. A promotion request requires the specification of the following two items:

- an object to promote
- a promotion policy governing the promotion process

A promotion policy identifies the rules used for a particular promotion. The policy determines which related repository objects, if any, are promoted along with the specified object. For example, when promoting a job, you often want to include any files referenced by the job steps. In addition, you may want to include execution server and credential definitions. The policy dictates how these items are handled.

Promotion is typically used in conjunction with label event notifications. For example, an analyst can set up the *Production* label to notify an administrator when the label is applied to a job so that the job can subsequently be promoted to production status.

Promotion re-creates the directory structure of the source repository server in the target repository server at the root (Content Repository folder) level. This allows location-based references between promoted objects to resolve correctly in the target server.

To submit a promotion request, the role for the user must include the *Promote Objects* action. In addition, the creation of new items in the target server may require that the role associated with the target credentials include certain actions. See the topic “Promotion considerations” for more information.

## Promoting objects

To promote an object:

1. In the Content Explorer, right-click on an object and select **Promote**. The Promote Object dialog box appears.
2. Specify the following information:
  - **Promote Version**. The version of the object to be promoted.
  - **Promote Policy**. The policy to be used for promotion.
3. If selected policy is for a delayed promotion, Export Location dialog box is displayed. Click **Browse** and specify the path of the export file.
4. Click **OK**. Depending on the specified policy, the object will be either promoted directly to the target repository or exported. If delayed promotion is performed, you must subsequently import the export file into the target repository. See the topic “Importing folders” on page 65 for more information.

## Promotion considerations

For an object, resource definitions, and referenced files to all be promoted successfully, the role for the target server credentials must include all relevant actions for the items included in the promotion request.

The following table identifies the actions needed for a variety of items that may be included in the promotion set.

*Table 8. Actions affecting promotion.*

Included in promotion set	Action needed for target credentials
Resource having notifications	<i>Define and Manage Notifications</i>
Credentials	<i>Define Credentials</i>
Resource having custom properties	<i>Define Custom Properties</i>

Table 8. Actions affecting promotion (continued).

Included in promotion set	Action needed for target credentials
Data source	<i>Define Datasources</i>
Message domain	<i>Define Message Domains</i>
Server cluster definition	<i>Define Server Clusters</i>
Server definition	<i>Define Servers</i>
Resource having topics	<i>Define Topics</i>
Job	<i>Job Edit</i>
Enterprise View, Application View, or Data Provider Definition	<i>Manage Enterprise View</i>
Resource having subscriptions for a different principal	<i>Create Subscriptions</i>
Resource having subscriptions for the target principal	<i>Manage Subscriptions</i>
Job having a schedule	<i>Schedules</i>
Resource version that is not the latest version in the target server	<i>Show All Versions</i>
Latest resource version	<i>Show LATEST</i>

In addition, the target credentials must have the *Manage label* permission for any label associated with a promoted object version.

---

## Chapter 8. Analytic data views

An analytic data view defines a structure for accessing data that describes the entities used in predictive models and business rules. The view associates the data structure with physical data sources for the analysis.

Predictive analytics requires data organized in tables with each row corresponding to an entity for which predictions are made. Each column in a table represents a measurable attribute of the entity. Some attributes may be derived by aggregating over the values for another attribute. For example, the rows of a table could represent customers with columns corresponding to the customer name, gender, zip code, and the number of times the customer had a purchase over \$500 in the past year. The last column is derived from the customer order history, which is typically stored in one or more related tables.

The predictive analytic process involves using different sets of data throughout the lifecycle of a model. During initial development of a predictive model, you use historic data that often has known outcomes for the event being predicted. To evaluate the model effectiveness and accuracy, you validate a candidate model against different data. After validating the model, you deploy it into production use to generate scores for multiple entities in a batch process or for single entities in a real-time process. If you combine the model with business rules in a decision management process, you use simulated data to validate the results of the combination. However, although the data that is used differs across the model development process stages, each data set must provide the same set of attributes for the model. The attribute set remains constant; the data records being analyzed change.

An analytic data view consists of the following components that address the specialized needs of predictive analytics:

- A data view schema, or data model, that defines a logical interface for accessing data as a set of attributes organized into related tables. Attributes in the model can be derived from other attributes.
- One or more data access plans that provide the data model attributes with physical values. You control the data available to the data model by specifying which data access plan is active for a particular application.

### **Important:**

- Components of the analytic data view are defined by using IBM SPSS Modeler streams. Familiarity with IBM SPSS Modeler concepts and experience creating streams are necessary to work with analytic data views.
- Adapters for IBM SPSS Modeler must be installed for your IBM SPSS Collaboration and Deployment Services Repository to define analytic data view components. For more information about these adapters, see the IBM SPSS Modeler documentation.




Figure 2. Analytic data view

Figure 2 illustrates an analytic data view that contains two data access plans for a data model. The data model includes three tables, with relationships defined between tables 1 and 2 and between tables 2 and 3. Data access plan 1 associates each table with an IBM SPSS Modeler stream. Data access plan 2 associates the data model tables with three different streams. Under data access plan 1, the model retrieves data from terminal nodes in stream 11, stream 12, and stream 13. Under data access plan 2, the model retrieves data from terminal nodes in stream 21, stream 22, and stream 23. By changing the data access plan in use, you can switch the data available to the model.

## Creating analytic data views

By creating an analytic data view, you define a container for a data model and its associated data access plans. This structure provides a user-defined interface for accessing data.

### Procedure

1. In the **Content Explorer**, select the folder in which to save the analytic data view.
2. From the menus, select **File > New > Analytic Data View**.
3. Enter a name for the data view.
4. Optional: Select **Lock** to prevent other users from modifying the object settings. If you lock the object, you are the only user able to edit the analytic data view until you unlock the object.
5. Click **Finish**.

### Results

The analytic data view is created in the selected folder and the analytic data view editor opens.

### What to do next

Create a data model and data access plan for the analytic data view.

---

## Data access plans

A data access plan associates the data model tables in an analytic data view with physical data sources. A data source corresponds to a terminal node from an IBM SPSS Modeler stream.

To be used as a source for a data model table, the data structure for the terminal node must correspond to the table structure. The measurement types for the terminal node fields must match the types for the data model tables. For example, if the data model table includes a field for gender that expects values of 0 and 1, the terminal node field that is mapped to this table field must be of the integer type.

A stream that is used as a data source can include most IBM SPSS Modeler nodes before the terminal node. For example, the stream can include multiple source nodes and use a **Merge** node to combine the data. You can use **Field Operations** nodes to create specific data structures. However, streams that include Enterprise View nodes are not supported.

A data access plan can use the source nodes as they are defined in the stream, or you can override them with alternative settings. The type of the data access plan determines how the source nodes are replaced. The following types are available:

- A batch data access plan specifies new values for source node parameters in the stream. For example, you can specify a different file for a **Variable File** node or a different table for a **Database** node.

**Restriction:** A batch plan cannot replace a source node with a source node of another type. For example, you cannot replace a **Database** node with a **Variable File** node.

- A real-time data access plan replaces source nodes with data source definitions stored in the IBM SPSS Collaboration and Deployment Services Repository. For example, you can replace a source node with an application server data source.

An analytic data view typically contains multiple data access plans. You specify the plan to use when you access the analytic data view in your application. For example, you could use a batch plan when training a model and a real-time plan when generating scores based on the trained model. The same stream is used in both instances, but the two plans access different data. You could also create data access plans for an analytic data view that are based on entirely different streams, or on different terminal nodes in the same stream.




Figure 3. Data access plans for a data source

Figure 3 illustrates two different data access plans for a stream that is used as a data source in an analytic data view. The stream reads data by using two source nodes, one for a database and one for the IBM SPSS Statistics data file `data.sav`. A series of subsequent nodes merges the data, restructures the content, and derives a new field. The terminal node presents the results as a table.

Two different data access plans specify settings for the source nodes in the stream. The batch plan uses the database settings as defined in the stream for the database node. For the data file, this plan substitutes a new data file, `testdata.sav`. In contrast, the real-time data access plan replaces the database node with a JDBC data source, and replaces the data file node with a data service data source.

Regardless of which data access plan is in use, the data structure that is defined by the stream must be satisfied. The `testdata.sav` file and the data service data source must both provide fields of the same type as the original `data.sav` file in the stream. Similarly, the JDBC data source must provide fields of the same type as the original database node. By maintaining the fields and their types, you ensure that the merge, restructure, and derive operations can successfully operate on the data to produce a table that contains the required fields of the correct type for an associated data model table.

**Important:** If a terminal node associated with an analytic data view table becomes unavailable, the mapping of the table to physical data is removed from any data access plans that use the terminal node as a data source. For example, if the version of the stream containing the terminal node expires or is removed from the IBM SPSS Collaboration and Deployment Services Repository, any tables associated with the stream need to be associated with a different data source before you can use the analytic data view.

## Creating data access plans

Associate analytic data view tables with physical data sources by using data access plans.



## Before you begin

- Store any IBM SPSS Modeler streams to be used as data sources for analytic data view tables in the IBM SPSS Collaboration and Deployment Services Repository.
- From the **Content Explorer**, open the analytic data view to contain the data access plan .

## About this task

An analytic data view typically contains multiple data access plans. You can change the data that is being accessed for the analytic data view tables by selecting a different data access plan for the view.

## Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. In the **Plan** field, click the down arrow and select **<Create new data access plan>**.
3. In the New Data Access Plan dialog box, enter a unique name for the data access plan in the **Name** field.
4. Select the plan type.
  - A **batch** data access plan uses the source nodes in the streams to access the data. You can use the node settings as specified in the streams or override the settings with new values. For example, if a stream contains a **Statistics File** node that references the file `data.sav`, you can change the referenced file to `testdata.sav` in the data access plan.
  - A **real-time** data access plan allows values to be specified interactively. You can override the source nodes in the data source streams with data providers to access the data. For example, if a stream contains a **Statistics File** node that references the file `data.sav`, you can override the entire node with a JDBC data source in the data access plan.
5. Click **OK**.
6. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The analytic data view includes a data access plan having the specified name and type.

## What to do next

- Add one or more tables to the analytic data view.
- Modify the data mapped to the analytic data view tables by the data access plan.
- Override the source nodes in the mapped IBM SPSS Modeler streams to access different data.

## Adding mapped tables to analytic data views

To add a mapped table to an analytic data view, you select the terminal node of an IBM SPSS Modeler stream on which to base the table. The table attributes are generated from the fields available in the terminal node.

## Before you begin

- Store any IBM SPSS Modeler streams to be used as data sources for analytic data view tables in the IBM SPSS Collaboration and Deployment Services Repository.
- From the **Content Explorer**, open an analytic data view.

## About this task

An analytic data view table corresponds to the terminal node of an IBM SPSS Modeler stream. The node fields determine the attributes and attribute types included in the table.

For example, suppose that the terminal node of a stream has the following fields and storage types:

- CustomerID as an integer
- Gender as a string
- Income as a real

A table that is based on this terminal node has the attributes CustomerID, Gender, and Income. The attribute types correspond to the storage type as defined in the stream.

## Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. Select the data access plan for the mapping of data to the new table.
3. In the Data View Table panel, click **New**.
4. In the New Data View Table dialog box, click **Browse** and select the IBM SPSS Modeler stream from the IBM SPSS Collaboration and Deployment Services Repository that defines the fields to be mapped to the analytic data view table attributes.
5. Select the label that specifies the version of the selected stream to be used.
6. Select the terminal node that specifies the fields to be used for the data model table.
7. Click **OK**. A table with a name corresponding to the terminal node is added to the analytic data view. If the name is not unique, it is automatically modified to prevent conflicts with other tables.
8. Specify the amount of time (in seconds) to cache the table during data retrieval. Caching tables can improve overall performance by preventing retrieval of the same data repeatedly.
9. Optional: Right-click the table name in the analytic data view editor and select **Rename** to modify the table name.
10. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The data model for the analytic data view includes the new table. The table contains attributes that correspond to the fields in the specified IBM SPSS Modeler stream. The physical data that is mapped to the table corresponds to the data from the stream. This data mapping is part of the current data access plan

## What to do next

- Add more tables to the analytic data view.
- Define the relationships that exist between the tables in the data model.
- Add derived attributes to the tables.
- Create a data access plan to define other data mappings for the analytic data view tables.
- When the analytic data view contains all needed tables and data access plans, use the view in your application to access the data. For example, you can use the analytic data view as a data provider for scoring configurations. You can also reference the analytic data view by using a Data View node in an IBM SPSS Modeler stream.

## Modifying the data mapping for analytic data view tables

You can change the IBM SPSS Modeler stream terminal node mapped to an analytic data view table to modify the data associated with the table or to update the overall table structure.

### Before you begin

- Store any IBM SPSS Modeler streams to be used as data sources for analytic data view tables in the IBM SPSS Collaboration and Deployment Services Repository.
- From the **Content Explorer**, open an analytic data view that contains at least one table.

### About this task

When you add a mapped table to an analytic data view, the current data access plan automatically associates the table attributes with the data from the selected terminal node. You can change the mapped stream to associate the table with different data. By changing the mapped stream, you can also change the overall data model structure. For example, you can add new attributes to existing tables.

If the analytic data view contains a table that is not mapped to any physical data, you must map the table to a stream terminal node before the analytic data view can be used to access data for attributes of that table.

### Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. Select the data access plan for the new mapping of data to the table.
3. Select the table to be mapped to a different terminal node.
4. Click **Map**.
5. In the Map Stream to Data View Table dialog box, click **Browse** and select the IBM SPSS Modeler stream from the IBM SPSS Collaboration and Deployment Services Repository that defines the fields to be mapped to the analytic data view table attributes.
6. Select the label that specifies the version of the selected stream to be used.
7. Select the terminal node that specifies the fields to be used for the data model table.
8. Optional: Click **Clear** to remove the current stream, label, and terminal node settings. The table is not mapped to any physical data source if you clear these settings. You will need to map the table to a data source before you can use it in any analyses.
9. Click **OK**.
10. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

### Results

The table attributes are mapped to fields from the selected terminal node that have the same name. If an attribute type does not match the type for the new mapped field, the Data View Schema displays a warning icon next to the attribute name to indicate that the mapping requires additional attention. Table attributes that do not correspond to fields in the new terminal node are not automatically mapped to any data in the current data access plan and need to be mapped manually.

### What to do next

After mapping a table to a new terminal node, you should review the automatic attribute mappings and modify any as needed manually.

## Mapping table attributes to stream fields

Manually map table attributes to the fields available in an IBM SPSS Modeler stream terminal node to manage the data that is associated with the attributes.

### Before you begin

From the **Content Explorer**, open an analytic data view that contains at least one table.

### About this task

When you map an analytic data view table to an IBM SPSS Modeler stream terminal node, the table attributes are automatically mapped to the node fields by using the attribute and field names. You can manually remap table attributes to associate the attributes with different data.

### Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. Select the data access plan.
3. Select the table containing the attributes to be mapped.
4. Expand the Field Map section to view the current attribute mapping.
5. Click the **Data View Attributes** cell for the stream field to be mapped. The cell displays a list that includes a blank value, **<Create new attribute>**, and all unmapped data view attributes in the analytic data view.
6. Specify the mapping for the stream field.
  - Select an unmapped attribute to map to the stream field.
  - Select the blank value to unmap the stream field.
  - Select **<Create new attribute>** to create a table attribute for the field. Specify a unique name for the attribute in the New Data View Attribute dialog box. The attribute type corresponds to the field type.
7. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

### Results

The table attributes are mapped to the stream fields as defined in the Field Map.

### What to do next

After mapping attributes to different fields, preview the data to verify that the correct values are being accessed.

## Overriding batch data access plan data sources

You can specify settings for IBM SPSS Modeler stream source nodes that differ from the settings in the stream. The data that is accessed by the plan corresponds to the new source node settings.

### Before you begin

From the **Content Explorer**, open an analytic data view that includes a batch data access plan.

## Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. Select the batch data access plan that contains the data source settings to be overridden.
3. Select the data model table to be associated with a different data source.
4. In the **Override Data Sources** section, select the source node to be overridden. The Source Nodes list contains every source node included in the data source stream.
5. Specify new parameters for the selected source node. The parameters available depend on the source node type.
6. Optional: Click **Preview** to view the data that is accessed for the specified parameters.
  - a. In the Modeler Information dialog box, select the execution server to use to process the data source. The list of available servers includes all IBM SPSS Modeler servers that are defined as resources for the current IBM SPSS Collaboration and Deployment Services Repository.
  - b. Select the credential to use for accessing the server. The list of available credentials includes all credentials that are defined as resources for the current IBM SPSS Collaboration and Deployment Services Repository.
  - c. Click **Get Data**.
  - d. In the Preview Result dialog box, review the data to identify any issues.
7. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The data access plan contains the updated source node settings.

## What to do next

Verify the data access by previewing the data that results from the current plan settings.

## Overriding real-time data access plan data sources

For real-time data access plans, you can override the IBM SPSS Modeler stream source nodes with data providers. The data that is accessed by the plan corresponds to the data providers instead of the source node settings in the stream.

## Before you begin

From the **Content Explorer**, open an analytic data view that includes a real-time data access plan.

## About this task

**Important:** The fields in the data provider must match the fields defined in the data model. For example, if the data model includes a field named gender that contains string values, the data provider must provide string values for gender.

## Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. Select the real-time data access plan that contains the data source node to be overridden.
3. Select the data model table to be associated with a different data source.
4. In the **Override Data Sources** section, select the source node to be overridden. The Source Nodes list contains every source node included in the data source stream.

5. Select the data provider type and specify the data provider settings for the selected source node. The settings available depend on the data provider type.

Option	Description
<b>Context Data</b>	For a context data provider, the data set maps to a single table. The values in the table are provided in real-time. For example, when a score for a customer is based on credit rating and geocode, the credit score and geocode values are provided as the context data for the request. For this provider, specify the table to use as context data .
<b>JDBC Data Source</b>	For a JDBC data source provider, the data set maps to a single table or view within a database that is accessed by using JDBC. The following settings must be specified: <ul style="list-style-type: none"> <li>• <b>JDBC Connection.</b> Select from the list of JDBC data sources defined in the system. To create a new JDBC data source, click <b>New</b>.</li> <li>• <b>Credential.</b> Select the credential for accessing the JDBC data source from the list of credentials defined in the system. To create a new credential, click <b>New</b>.</li> <li>• <b>Table.</b> Select this option to specify a table from the data provider.</li> <li>• <b>Query.</b> Select this option to specify a SQL query that extracts data from the data provider. Either type the SQL statements in the <b>Edit query</b> field or click <b>Load From File</b> to select a text file containing the SQL statements.</li> </ul>
<b>Application Server Data Source</b>	For an application server data source provider, the data set maps to a single SQL table or view within a SQL database. You do not specify a credential when connecting to the data source; the credentials are defined in the application server. The following settings must be specified: <ul style="list-style-type: none"> <li>• <b>JNDI Connection.</b> Select from the list of application server data sources defined in the system. To create a new application server data source, click <b>New</b>.</li> <li>• <b>Table.</b> Select this option to specify a table from the data provider.</li> <li>• <b>Query.</b> Select this option to specify a SQL query that extracts data from the data provider. Either type the SQL statements in the <b>Edit query</b> field or click <b>Load From File</b> to select a text file containing the SQL statements.</li> </ul>
<b>Data Service Data Source</b>	For a data services data source provider, the data set maps to a single table that is defined in the data service. Select the source from the list of data service data sources defined in the system. To create a new data service data source, click <b>New</b> .

6. To limit the records being retrieved to those meeting specific criteria, select **Filters** to define filter conditions.
- a. Click **Add**.
  - b. In the Filter Information dialog box, select the column to use for filtering.
  - c. Select the filter type.
  - d. Specify the table and value to use as the filter criterion.

7. For **Record Count Limit**, specify the maximum number of records to retrieve from the data source. If this value is exceeded during data retrieval, an error occurs. If you do not specify a value, the data access plan uses the default configuration value that is defined for your system. For more information, see the Data Service configuration settings available in the browser-based IBM SPSS Collaboration and Deployment Services Deployment Manager.
8. Optional: Click **Preview** and specify a value for the filter to view the data that is accessed for the specified data provider.
9. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The data access plan contains the updated data source settings.

## What to do next

Verify the data access by previewing the data that results from the current plan settings.

## Previewing data access plan data

You can verify that the correct data is being retrieved by a data access plan by previewing the data.

## Before you begin

From the **Content Explorer**, open an analytic data view that includes a data access plan.

## Procedure

1. Select the data access plan that contains the data source settings to be previewed.
2. Select the data view schema table to be previewed.
3. In the **Preview Data Sources** section, click **Preview**.
  - a. In the Modeler Information dialog box, select the execution server to use to process the data source. The list of available servers includes all IBM SPSS Modeler servers that are defined as resources for the current IBM SPSS Collaboration and Deployment Services Repository.
  - b. Select the credential to use for accessing the server. The list of available credentials includes all credentials that are defined as resources for the current IBM SPSS Collaboration and Deployment Services Repository.
  - c. Click **Get Data**.
  - d. In the Preview Result dialog box, review the data to identify any issues.

## Results

Several values for the data model table fields are displayed.

## Deleting data access plans

Remove a data access plan from an analytic data view if the set of stream and field mappings are no longer needed.

## Before you begin

From the **Content Explorer**, open the analytic data view that contains the data access plan to be deleted.

## Procedure

1. In the analytic data view editor, select the **Data Access Plans** tab.
2. In the **Plan** field, click the down arrow and select the data access plan to be deleted as the current data access plan

**Note:** An analytic data view must contain at least one data access plan. If the view contains only one access plan, you cannot delete the plan.

3. Click **Delete**.
4. In the warning message dialog box, click **OK**.

## Results

The data access plan is removed from the analytic data view.

---

## Data models

The data model for an analytic data view is a schema that defines the interface through which data is processed as a set of one or more tables. Each table in the data model represents a concept or entity involved in the predictive analytic process. Attributes for the tables correspond to attributes of the entities represented by the tables.

For example, if you are analyzing customer orders, your data model could include a table for customers and a table for orders. The customers table might have attributes for customer identifier, age, gender, marital status, and country of residence. The orders table might have attributes for the order identifier, number of items in the order, total cost, and the identifier for the customer who placed the order. The customer identifier attribute could be used to associate the customers in the customers table with their orders in the orders table.

Table attributes can be derived from attributes in related tables. For example, if you need the number of orders for each customer, you can include an order count attribute in the customers table. The value of this attribute is derived by aggregating over the orders table, counting the number of orders for each customer.

## Relationships

A data model includes relationships between the model tables that describe how the entities represented by the tables are related. The relationship cardinality indicates how the rows of a table relate to the rows of another table, and determines how the data in the tables is combined.

The cardinality of a relationship between two tables is defined as one of the following types:

- *One-to-one*. One row of table A corresponds to at most one row of table B. Each row of table B corresponds to at most one row of table A.
- *One-to-many*. One row of table A corresponds to any number of rows of table B. Each row of table B corresponds to at most one row of table A.
- *Many-to-many*. One row of table A corresponds to multiple rows of table B. Each row of table B corresponds to multiple rows of table A.



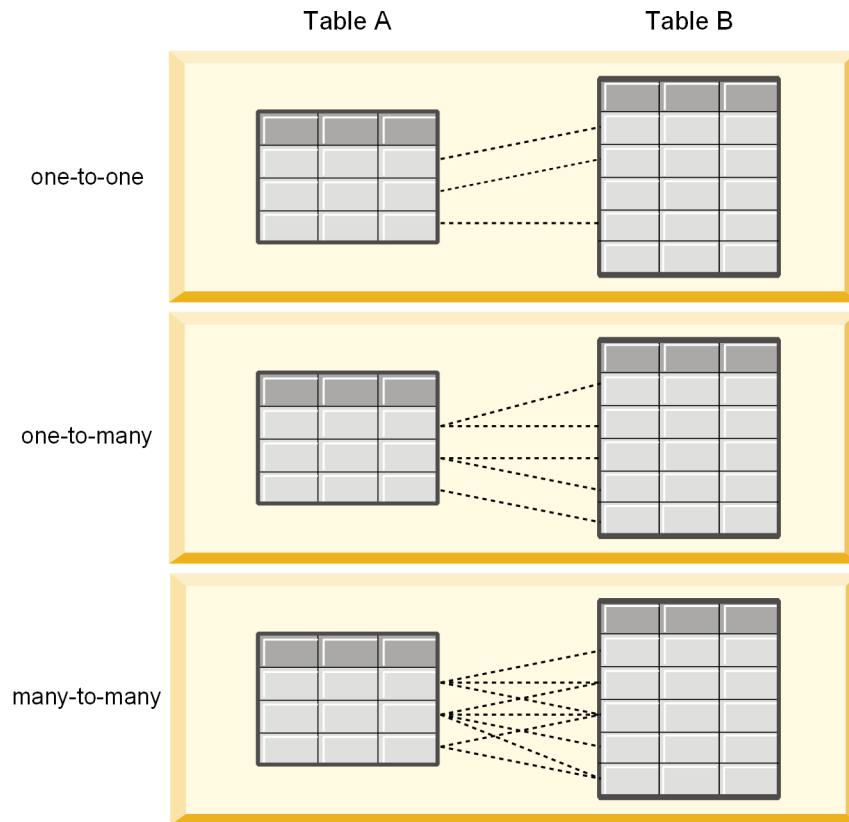


Figure 4. Table relationship types

Figure 4 illustrates the three relationship types. In the one-to-one relationship, each row of Table A is connected to one row of Table B. For example, suppose Table A represents reward program membership and Table B represents customers. A membership identifier in Table A is only associated with a single customer who participates in the program. Some of the customers in Table A do not participate in the rewards program, but the participating customers are associated with one membership.

In the one-to-many relationship, the first row of Table A is connected to two rows of Table B. The second row of Table A is also connected to two rows of Table B. The third row of Table A is connected to the last row of Table B. For Table B, each row is connected to only one row of Table A. For example, suppose Table A represents customers and Table B represents orders. Each customer in Table A can make multiple orders. However, each order in Table B can be associated only with the one customer who placed the order.

In the many-to-many relationship, The first row of Table A is connected to three rows of Table B. The second row of Table A is connected to four rows of Table B. The third row of Table A is connected to two rows of Table B. For Table B, the first row is connected to one row of Table A. The second row of Table B, however, is connected to two rows of Table A. Any row in either table can be connected to any number of rows in the other table. For example, suppose Table A represents orders and Table B represents products. Any order in Table A can include any number of products. Furthermore, any product in Table B can be included in multiple orders.

**Note:** All table relationships in an analytic data view have a cardinality of one-to-many.

## Specifying the analytic data view editor advanced mode

To modify the data model for an analytic data view, you must enable the advanced mode.

## About this task

Many analyses require an analytic data view that contains a data model that consists of a single table. If your data model requires a more complex structure, you must enable the advanced mode. In this mode, you can manually define tables and attributes, create derived attributes, and specify table relationships. You can import a data model corresponding to an IBM Operational Decision Manager Business Object Model (BOM), or export a data model as an archive file.

### Procedure

1. In the **Content Explorer**, right-click the server name and select **Properties**.
2. In the Properties dialog box, open the **User Preferences** section and select **Analytic Data View Editor Advanced Mode**.
3. Select **Advanced Mode** to enable editing of the data model for analytic data views. Clear this option to disable data model editing.
4. Click **OK**.

### Results

If you enable the advanced mode, the **Data View Configuration** tab is available when you open the analytic data view editor. If you disable this mode, the tab is not available.

## Adding unmapped tables to analytic data views

You can add tables to analytic data views manually if they are needed by applications that access the data view.

### Before you begin

- Enable the advanced mode for the analytic data view editor.
- From the **Content Explorer**, open an analytic data view for editing.

## About this task

When you define mapped tables, the data model adds attributes to the tables based on the fields available in the associated stream. Alternatively, you can add unmapped tables manually.

### Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Right-click the area under the Data View Schema and select **New > New Table**.
3. Right-click the new table and select **Rename** to enter a name for the table.  
The following steps involve properties related to business object models. If you do not plan on exporting the data model to a BOM archive, the default values are sufficient.
4. In the Properties panel, define the table verbalization. The table verbalization determines the phrases available when creating business rules based on the table.
5. Optional: Select **Static** to make the table available at the class level in the object model.
6. Optional: Select **Final** to ensure that the table cannot be changed in the object model.
7. Optional: Select **Deprecated** to mark the table as deprecated.
8. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The data model table includes the new table.

## What to do next

After adding an unmapped table, add attributes to the table.

## Adding unmapped attributes to data model tables

You can add attributes to data model tables manually if they are needed by applications that access the data view.

### Before you begin

- Enable the advanced mode for the analytic data view editor.
- From the **Content Explorer**, open an analytic data view for editing.

### About this task

When you define mapped tables, the data model adds attributes to the tables based on the fields available in the associated stream. Alternatively, you can add fields to the data model tables manually.

### Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Select a table in the data view schema.
3. Right-click and select **New > New Attribute**.
4. Right-click the new attribute and select **Rename** to enter a name for the attribute.
5. In the Properties panel, specify the type of data associated with the attribute.
  - **String**. Sequences of alphanumeric characters.
  - **Integer**. Whole numbers.
  - **Real**. Fractional numeric values.
  - **Date**. Values representing calendar dates.
  - **Time**. Values representing time points.
  - **Timestamp**. Values containing both date and time information.

The following steps involve properties related to business object models. If you do not plan on exporting the data model to a BOM archive, the default values are sufficient.

6. Specify the read/write properties for the field. The read/write property determines the methods included in the object model and the verbalization associated with the attribute.
  - **Read/Write**. The object model includes a get and a set method for the attribute.
  - **Read Only**. The object model includes only a get method for the attribute.
  - **Write Only**. The object model includes only a set method for the attribute.
7. Define the attribute verbalization. The attribute verbalization determines the phrases available when creating business rules based on the attribute.
  - **Read**. The verbalization specifies a navigation phrase for the attribute. The structure of a navigation phrase is typically *{attribute\_name} of {this}*, where *{this}* represents the entity in the object model described by the attribute.
  - **Write**. The verbalization specifies an action phrase for the attribute. The structure of an action phrase is typically *set the {attribute\_name} of {this} to {attribute\_name}*, where *{this}* represents the entity in the object model described by the attribute.
8. Optional: Select **Static** to make the attribute available at the class level in the object model.

- Optional: Select **Final** to ensure that the attribute cannot be changed in the object model.
- Optional: Select **Deprecated** to mark the attribute as deprecated.
- From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The data model table includes the new attribute.

## What to do next

After adding attributes, map them to physical data by using a data access plan.

## Defining relationships between analytic data view tables

Specify how the data records in multiple tables are combined by defining relationships between the tables in an analytic data view.

### Before you begin

- From the **Content Explorer**, open an analytic data view that contains at least two tables.
- Enable the advanced mode for the analytic data view editor.

### About this task

You define a relationship between tables by creating a collection attribute. This attribute specifies a table that contains data on which attributes of the original table can be based.

### Procedure

- In the analytic data view editor, select the **Data View Configuration** tab.
- Select a table in the data model.
- Click **Add a new collection attribute to the selected table**.
- In the Properties pane, enter a name for the collection attribute.
- For the type of collection, select the other data model table to include in the relationship with the current table.
- Optional: If you plan on exporting the data model to a BOM archive, specify the read/write and verbalization properties. For information about specific properties, see "Modifying data model properties" on page 88
- In the Relationship section, select an attribute in the **From Table** and an attribute in the **To Table** to be linked in the relationship. The two selected attributes must be of the same type.
- Click **Map**.
- Optional: To remove a mapping between two attributes, select one of the attributes and click **Unmap**.
- Optional: To remove all mappings between attributes, click **Unmap All**.
- From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

One table contains a collection attribute that references a second table.

### What to do next

- Create a derived attribute based on the relationship
- Define relationships for other table pairs in the data model

## Adding derived attributes to analytic data view tables

You can add attributes to analytic data view tables that are computed from the values of other table attributes.

### Before you begin

- Enable the advanced mode for the analytic data view editor.
- In an analytic data view, create a relationship between the table that will contain the derived attribute and the table containing the attributes on which the derived attribute will be based.

### About this task

A derived attribute for a table has values that are computed from attributes in related tables. The definition of a derived attribute has the following structure:

**<aggregation expression>** where **<condition expression>**

The parameters **<aggregation expression>** and **<condition expression>** correspond to verbalization elements from the data model, aggregation operators, and conditional operators.

### Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Select a table in the data model.
3. Click **Add a new derived attribute to the selected table**.
4. In the Properties pane, enter a name for the field.
5. In the Properties panel, specify the type of data associated with the attribute.
  - **String**. Sequences of alphanumeric characters.
  - **Integer**. Whole numbers.
  - **Real**. Fractional numeric values.
  - **Date**. Values representing calendar dates.
  - **Time**. Values representing time points.
  - **Timestamp**. Values containing both date and time information.
6. Define the attribute verbalization. The verbalization determines the phrases available when creating business rules based on the attribute. The verbalization is also used as the description of the attribute in some applications that use the analytic data view, such as IBM SPSS Modeler.
7. Specify the definition for the derived attribute. Either type the complete expression in the **Definition** field or press the spacebar to interactively build the expression by using context assist.
  - a. Press the spacebar and select **<aggr>**.
  - b. Select the aggregation function and specify the function arguments.
  - c. Optional: To include a conditional expression after the aggregation expression in the definition, select where **<condition>** and construct the conditional expression using the operators and verbalization elements.
8. Optional: Select **Static** to make the table or attribute available at the class level in the object model.
9. Optional: Select **Final** to ensure that the table or attribute cannot be changed in the object model.

10. Optional: Select **Deprecated** to mark the table or attribute as deprecated.
11. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The analytic data view table includes the new derived field.

## Example

Suppose a data model includes a table that contains customer information and a table that contains order information. A derived attribute that indicates the order total for each customer for the third quarter could have the following definition:

**the total amount of the orders of this customer**  
where **the order date of each order is after 7/1/2013 and before 9/30/2013**

## What to do next

Test the derived attribute by previewing the data associated with the table.

## Modifying data model properties

You can export an analytic data view for use in IBM Operational Decision Manager. Properties of the analytic data view determine how the model is processed.

## Before you begin

- Enable the advanced mode for the analytic data view editor.
- From the **Content Explorer**, open an analytic data view for editing.

## About this task

The analytic data view includes metadata that describes the tables and attributes in the model. The default property values will suffice for many applications. However, if you will be using the analytic data view tables in IBM Operational Decision Manager, you can modify the metadata to control how the model is processed.

## Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Select either a table or an attribute in the data model and specify its properties. The properties available depend on your selection.

Option	Description
<b>Table</b>	For a table, specify the following properties: <ul style="list-style-type: none"> <li>• <b>Name.</b> The name of the table, which corresponds to the class name in the object model.</li> <li>• <b>Superclass.</b> The class in the object model from which the current class is derived, if any.</li> <li>• <b>Interface.</b> The interface implemented by the class, if any.</li> <li>• <b>Verbalization.</b> The table verbalization determines the phrases available when creating business rules based on the table.</li> </ul>
<b>Attribute</b>	For an attribute, specify the following properties: <ul style="list-style-type: none"> <li>• <b>Name.</b> The name of the attribute in the object model.</li> <li>• <b>Type.</b> The type of data associated with the attribute.               <ul style="list-style-type: none"> <li>– <b>String.</b> Sequences of alphanumeric characters.</li> <li>– <b>Integer.</b> Whole numbers.</li> <li>– <b>Real.</b> Fractional numeric values.</li> <li>– <b>Date.</b> Values representing calendar dates.</li> <li>– <b>Time.</b> Values representing time points.</li> <li>– <b>Timestamp.</b> Values containing both date and time information.</li> </ul> </li> <li>• <b>Read/Write.</b> The read/write property determines the methods included in the object model and the verbalization associated with the attribute.               <ul style="list-style-type: none"> <li>– <b>Read/Write.</b> The object model includes a get and a set method for the attribute.</li> <li>– <b>Read Only.</b> The object model includes only a get method for the attribute.</li> <li>– <b>Write Only.</b> The object model includes only a set method for the attribute.</li> </ul> </li> <li>• <b>Verbalization.</b> The attribute verbalization determines the phrases available when creating business rules based on the attribute.               <ul style="list-style-type: none"> <li>– <b>Read.</b> The verbalization specifies a navigation phrase for the attribute. The structure of a navigation phrase is typically <i>{attribute_name}</i> of <i>{this}</i>, where <i>{this}</i> represents the entity in the object model described by the attribute.</li> <li>– <b>Write.</b> The verbalization specifies an action phrase for the attribute. The structure of an action phrase is typically set the <i>attribute_name</i> of <i>{this}</i> to <i>{attribute_name}</i>, where <i>{this}</i> represents the entity in the object model described by the attribute.</li> </ul> </li> </ul>

3. Optional: Select **Static** to make the table or attribute available at the class level in the object model.
4. Optional: Select **Final** to ensure that the table or attribute cannot be changed in the object model.
5. Optional: Select **Deprecated** to mark the table or attribute as deprecated.
6. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The data model elements have the new properties.

## What to do next

After you have specified the properties for all elements of the analytic data view, export the data model as a business object model or execution object model for subsequent use in IBM Operational Decision Manager.

## Deleting data model components

Remove tables or attributes from the data model if the items are no longer needed.

### Before you begin

- From the **Content Explorer**, open an analytic data view with at least one table or attribute to be deleted.
- Enable the advanced mode for the analytic data view editor.

### Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Select the item to be deleted in the data model.
3. Right-click and select **Delete**.
4. From the menus, select **File > Save** to save a version of the analytic data view that contains these settings.

**Note:** Every time that you save the analytic data view, a new version is created. To limit the number of versions that are created, make all necessary changes to the view before saving.

## Results

The table or attribute is no longer included in the data model.

---

## Business object models

A data model in an analytic data view has a one-to-one correspondence with a business object model (BOM) in IBM Operational Decision Manager. A BOM specifies the business elements and vocabulary used to define business rules. The vocabulary allows users to create business rules without any knowledge of how the data referenced in the rules is structured.

A business rule has the form "if **criteria** then **outcome**". For example, a decision on whether or not to reject a loan application based on the credit score of the applicant corresponds to the following business rule:

If the **credit score** of the **applicant** is less than 300 then reject the application.

The applicant is the business entity on which this rule is based. The credit score of the applicant is the entity attribute on which the decision is made. The business object model must include an element representing the applicant and that element must include an attribute corresponding to the credit score. The model also specifies the vocabulary used to reference the model elements in the rules. In this case the vocabulary includes "credit score of the applicant".

Structurally, the business object model is similar to the Java object model. Business entities correspond to a classes and can be grouped into packages. Classes can be nested in other classes. The entity attributes



correspond to class attributes that have a type that indicates the type of data values allowed for the attribute. For the example rule, the BOM includes an applicant class that has a credit score attribute of the integer type.

Elements of the data model for an analytic data view correspond to elements in the business object model. A data model table corresponds to a BOM class. A field in the data model table corresponds to an attribute for the class.

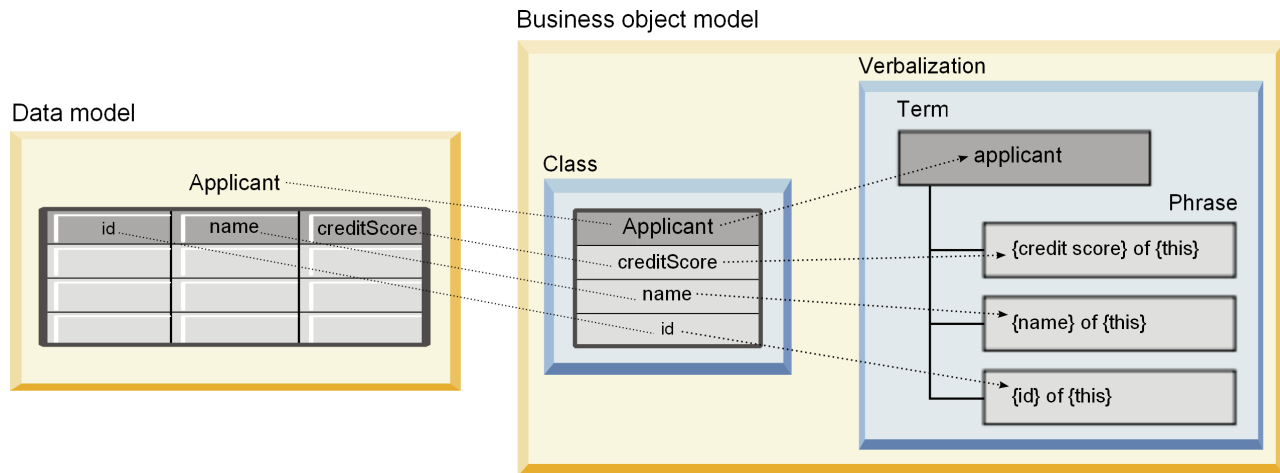


Figure 5. Relationships between the data model and the business object model

Figure 5 illustrates the relationships between the data model elements and the business object model elements for a data model that includes a table named Applicant. That table corresponds to a BOM class that has the same name. The verbalization of the class includes the term applicant for the class. The data model table fields of id, name, and creditScore correspond to attributes of the BOM class. Each attribute of the class yields a verbalization phrase of the form **{attribute}** of {this}. The variable {this} refers to the parent term applicant in the verbalization.

You can use the correspondence of the data model and the BOM to combine the results from different analytical approaches. The output from predictive models that are based on an analytic data view can be combined with business rules based on a BOM if the data model fields match the BOM elements. The common data interface allows scores from predictive models to be integrated with business rule criteria.

## Exporting data models as business object model archives

You can export a data model from an analytic data view as an IBM Operational Decision Manager business object model archive file. The data interface can then be used by any application that can process the business object model.

### Before you begin

- Enable the advanced mode for the analytic data view editor.
- From the **Content Explorer**, open an analytic data view that contains a data model with at least one table.

### About this task

The elements of a data model exported as a BOM archive can be used in IBM Operational Decision Manager to create business rules. Business rules that are based on this data interface use the same inputs as predictive models that are based on the data model. This correspondence facilitates combining predictive models and business rules.

A BOM archive file contains the following items:

- A .bom file that describes the business object model structure
- A .voc file that describes the verbalization for the business object model

**Note:** The archive file does not include a .b2x file that maps the business object model to an execution object model. You can create this mapping for your execution object model by using IBM Operational Decision Manager.

## Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Right-click the data model and select **Export > Export BOM**.
3. In the Save As dialog box, specify the name and location of the archive file to contain the exported BOM.
4. Click **Save**.

## Results

An archive that contains BOM artifacts corresponding to the data model is available at the specified location.

## What to do next

Use the archive in IBM Operational Decision Manager to create business rules.

## Importing business object models as data models

If the data model for your analytic data view is empty, you can populate the model with the structure defined in an IBM Operational Decision Manager business object model.

## Before you begin

- Enable the advanced mode for the analytic data view editor.
- From the **Content Explorer**, open an analytic data view containing an empty data model for editing.

## About this task

You can import a BOM archive file to create a data model that uses tables and fields that correspond to the elements in the BOM. Predictive models that are based on this data model use the same elements as business rules based on the BOM. This correspondence facilitates combining predictive models and business rules.

A BOM archive file contains the following items:

- A .bom file that describes the business object model structure
- A .voc file that describes the verbalization for the business object model

## Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Right-click the data model and select **Import**.
3. In the Open dialog box, select the archive that contains the BOM to be imported.
4. Click **Open**.

## Results

The data model includes the tables and fields that are defined in the BOM.

## What to do next

Associate the data model tables with IBM SPSS Modeler streams by using a data access plan.

---

## Execution object models

A data model in an analytic data view has a one-to-one correspondence with an execution object model (XOM) in IBM Operational Decision Manager. A XOM specifies the objects against which business rules are executed.

A XOM is derived from XML Schema Definitions. A XOM archive contains an .xsd schema file that contains definitions of the field types in the data model. Data model tables correspond to complexType entries. Each complexType entry contains a sequence of elements that correspond to the fields in that table. The type attribute for the elements designates the field type.

For example, suppose the data model for an analytic data view contains two tables. One table represents customers and contains fields for customer identifier, name, gender, birth date, zip code, a collection relationship, and the total of all orders. The other table represents orders and contains fields for order identifier, customer identifier, date, number of items, and order total. The schema file for this model follows:

```
<xsd:schema xmlns:xsd="http://www.w3.org/2001/XMLSchema"
attributeFormDefault="unqualified" elementFormDefault="unqualified"
targetNamespace="http://" xmlns="http://www.w3.org/2001/XMLSchema"
xmlns:xom="http://">
  <xsd:complexType name="Orders">
    <xsd:sequence>
      <xsd:element name="orderid" type="xom:int" />
      <xsd:element name="custid" type="xom:int" />
      <xsd:element name="date" type="xom:java.sql.Date" />
      <xsd:element name="numitems" type="xom:int" />
      <xsd:element name="total" type="xsd:double" />
    </xsd:sequence>
  </xsd:complexType>
  <xsd:complexType name="Customers">
    <xsd:sequence>
      <xsd:element name="custid" type="xom:int" />
      <xsd:element name="name" type="xsd:string" />
      <xsd:element name="gender" type="xsd:string" />
      <xsd:element name="birthdate" type="xom:java.sql.Date" />
      <xsd:element name="zipcode" type="xom:int" />
      <xsd:element maxOccurs="unbounded" minOccurs="0"
        name="ordersCollection" type="xom:Orders" />
      <xsd:element name="ordertotal" type="xsd:double" />
    </xsd:sequence>
  </xsd:complexType>
</xsd:schema>
```

The attributes minOccurs and maxOccurs for the field elements denote the minimum and maximum number of possible occurrences for the fields. If the attributes are not specified, the default value for both attributes is 1 and the field is required by any business rules that are based on the .xsd file.

## Exporting data models as execution object model archives

You can export a data model from an analytic data view as an IBM Operational Decision Manager execution object model archive file. The data interface can then be used by any application that can process the execution object model.

### Before you begin

- Enable the advanced mode for the analytic data view editor.

- From the **Content Explorer**, open an analytic data view that contains a data model with at least one table.

## About this task

The elements of a data model exported as a XOM archive can be used in IBM Operational Decision Manager to execute business rules. Business rules that are based on this data interface use the same inputs as predictive models that are based on the data model. This correspondence facilitates combining predictive models and business rules.

## Procedure

1. In the analytic data view editor, select the **Data View Configuration** tab.
2. Right-click the data model and select **Export > Export XOM**.
3. In the Save As dialog box, specify the name and location of the archive file to contain the exported XOM.
4. Click **Save**.

## Results

An archive that contains XOM artifacts corresponding to the data model is available at the specified location.

## What to do next

Use the archive in IBM Operational Decision Manager to create business rules.

---

## Chapter 9. Scoring

Scoring is the process of generating real-time values by supplying predictive models with input data. A scoring model is any artifact that can be used to produce output values given input data, such as a PMML file from IBM SPSS Statistics.

In general, to use a model for generating scores:

1. Select a model to use for scoring from the IBM SPSS Collaboration and Deployment Services Repository.
2. Define a scoring configuration for the model.
3. Supply the configured model with data and generate scores.

The predictive model used for scoring can be defined using IBM SPSS Modeler streams, scenarios, or PMML generated from IBM Corp. products. IBM SPSS Collaboration and Deployment Services supports PMML 3.2 and earlier versions. In addition, legacy markup (e.g. SPSS-ML) from older products may also be used for scoring. Note that some nodes in streams, such as the ADP node, must be trained before they can be used for scoring. For more information, see the IBM SPSS Modeler documentation.

**Important:** To use a particular model type for scoring, a scoring adapter for that model type must be installed on the IBM SPSS Collaboration and Deployment Services Repository server. For example, to generate scores based on PMML files, the IBM SPSS Collaboration and Deployment Services Scoring Adapter for PMML must be installed. To use IBM SPSS Modeler files for scoring, the IBM SPSS Collaboration and Deployment Services Scoring Adapter for IBM SPSS Modeler must be installed.

Use the IBM SPSS Collaboration and Deployment Services Deployment Manager to define model scoring configurations and to monitor the model scoring performance. Results generated from scoring can be viewed in the IBM SPSS Collaboration and Deployment Services Deployment Portal or in custom client applications.

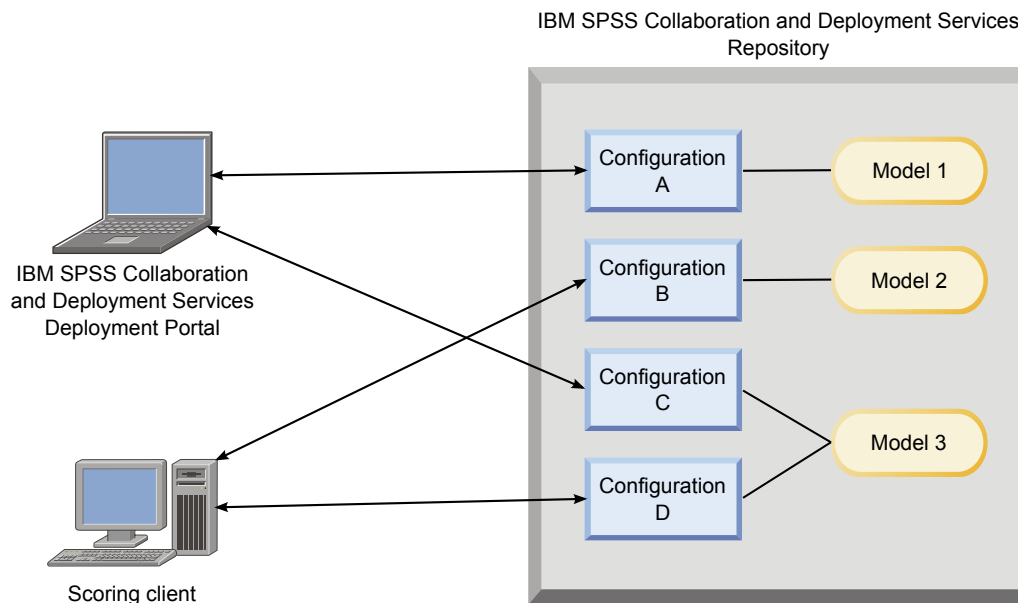


Figure 6. General scoring architecture

Figure 6 on page 95 illustrates the general architecture for scoring. In this example, the IBM SPSS Collaboration and Deployment Services Repository contains three predictive models. Configuration A defines the scoring settings for Model 1. Configuration B defines the scoring settings for Model 2. Configurations C and D define different scoring settings for Model 3.

IBM SPSS Collaboration and Deployment Services Deployment Portal users send separate scoring requests to models 1 and 3 that use the settings in configurations A and C. The scoring service generates the scores and returns the values to the users.

As an alternative to using IBM SPSS Collaboration and Deployment Services Deployment Portal, you can create a custom scoring client to send scoring requests and to receive results. In this example, a custom client sends separate scoring requests to models 2 and 3 that use the settings in configurations B and D. The scoring service generates the scores and returns the values to the client.

---

## Supported scoring functions and models

Applying a predictive model to a set of data can produce a variety of scores, such as predicted values, predicted probabilities, and other values based on that model. The type of score produced is referred to as the **scoring function**. The following scoring functions are available:

Scoring function	Description
<b>PREDICT</b>	Returns the predicted value of the target variable.
<b>STDDEV</b>	Standard deviation.
<b>PROBABILITY</b>	Probability associated with a particular category of a target variable. Applies only to categorical variables. In the absence of the optional third parameter, <i>category</i> , this is the probability that the predicted category is the correct one for the target variable. If a particular category is specified, then this is the probability that the specified category is the correct one for the target variable.
<b>CONFIDENCE</b>	A probability measure associated with the predicted value of a categorical target variable. Applies only to categorical variables.
<b>NODEID</b>	The terminal node number. Applies only to tree models.
<b>CUMHAZARD</b>	Cumulative hazard value. Applies only to Cox regression models.
<b>NEIGHBOR</b>	The ID of the <i>k</i> th nearest neighbor. Applies only to nearest neighbor models. In the absence of the optional third parameter, <i>k</i> , this is the ID of the nearest neighbor. The ID is the value of the case labels variable, if supplied, and otherwise the case number.
<b>DISTANCE</b>	The distance to the <i>k</i> th nearest neighbor. Applies only to nearest neighbor models. In the absence of the optional third parameter, <i>k</i> , this is the distance to the nearest neighbor. Depending on the model, either Euclidean or City Block distance will be used.

The following table lists the set of scoring functions available for each type of model that supports scoring. The function type denoted as PROBABILITY (category) refers to specification of a particular category (the optional third parameter) for the PROBABILITY function.

Table 9. Supported functions by model type.

Model type	Supported functions
Tree (categorical target)	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE, NODEID
Tree (scale target)	PREDICT, NODEID, STDDEV

Table 9. Supported functions by model type (continued).

Model type	Supported functions
Boosted Tree (C5.0)	PREDICT, CONFIDENCE
Linear Regression	PREDICT, STDDEV
Automatic Linear Models	PREDICT
Binary Logistic Regression	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Conditional Logistic Regression	PREDICT
Multinomial Logistic Regression	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
General Linear Model	PREDICT, STDDEV
Discriminant	PREDICT, PROBABILITY, PROBABILITY (category)
TwoStep Cluster	PREDICT
K-Means Cluster	PREDICT
Kohonen	PREDICT
Neural Net (categorical target)	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Neural Net (scale target)	PREDICT
Naive Bayes	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Anomaly Detection	PREDICT
Ruleset	PREDICT, CONFIDENCE
Generalized Linear Model (categorical target)	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Generalized Linear Model (scale target)	PREDICT, STDDEV
Generalized Linear Mixed Model (categorical target)	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Generalized Linear Mixed Model (scale target)	PREDICT
Ordinal Multinomial Regression	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Cox Regression	PREDICT, CUMHAZARD
Nearest Neighbor (scale target)	PREDICT, NEIGHBOR, NEIGHBOR(K), DISTANCE, DISTANCE(K)

Table 9. Supported functions by model type (continued).

Model type	Supported functions
Nearest Neighbor (categorical target)	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE, NEIGHBOR, NEIGHBOR(K), DISTANCE, DISTANCE(K)
Bayesian Network	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Support Vector Machine (categorical target)	PREDICT, PROBABILITY, PROBABILITY (category), CONFIDENCE
Support Vector Machine (scale target)	PREDICT, STDDEV

- For the Binary Logistic Regression, Multinomial Logistic Regression, and Naive Bayes models, the value returned by the CONFIDENCE function is identical to that returned by the PROBABILITY function.
- For the K-Means model, the value returned by the CONFIDENCE function is the least distance.
- For tree and ruleset models, the confidence can be interpreted as an adjusted probability of the predicted category and is always less than the value given by PROBABILITY. For these models, the confidence value is more reliable than the value given by PROBABILITY.
- For neural network models, the confidence provides a measure of whether the predicted category is much more likely than the second-best predicted category.
- For Ordinal Multinomial Regression and Generalized Linear Model, the PROBABILITY function is supported when the target variable is binary.
- For nearest neighbor models without a target variable, the available functions are NEIGHBOR and DISTANCE.

## Scoring configurations

Before a model can be used for scoring, supplemental information must be defined. For example, an IBM SPSS Modeler stream containing a Data View source node needs tables from a data access plan specified to retrieve required data.

This additional information constitutes a scoring configuration for the model, and defines scoring parameters such as:

- Identification information for the configuration itself
- Identification information for the model used for scoring
- The data provider for the input
- Settings for logging
- The order of the input attributes
- Cache size used for scoring models

A single model may be used in a variety of scoring situations that require different scoring parameters. For example, scores may be based on a test data provider for internal purposes and on a different data provider for production usage. Alternatively, the information being logged as result of scoring may depend on the scoring situation. To allow a model to be used in differing scoring circumstances, any model may be associated with multiple scoring configurations.

Scoring configurations can be suspended to temporarily prevent the processing of score requests. A suspended configuration must be reactivated before it can be used to generate scores.

To create a scoring configuration for a model, right-click the model to be configured for scoring in the Content Explorer and select **Configure Scoring**. The Scoring Configuration wizard appears.



**Important:** If the **Configure Scoring** option is not available for your model type, verify with your system administrator that the scoring adapter for that model type is installed for the IBM SPSS Collaboration and Deployment Services Repository server. When you connect to a repository server that includes scoring adapters, your environment updates to include the scoring functionality.

## New scoring model configuration

The Add New Scoring Model Configuration panel of the Scoring Configuration wizard identifies the model being used for scoring.

**Name.** The name for the scoring configuration. A model may have many scoring configurations, each identified by its unique name.

**Model File.** The name of the model file associated with the configuration.

**Label.** The label identifying the version of the model file configured for scoring.

Click **Next** to specify additional settings.

## Model specific settings

The Model Specific Settings panel of the Scoring Configuration wizard defines settings parameters for the specific model being configured. The settings parameters available vary across different models.

Specify values for the specific settings offered for the model. Click **Next** to specify additional settings.

## Data provider settings

To generate a score, you must provide the predictive model with values for the input fields. This data can be supplied manually by providing a value for each field or retrieved automatically by using a data source.

For example, suppose a model requires values for the input fields age, income, and gender. To generate a score for a customer, you can enter the value for that person for each field manually. Alternatively, you can use an identifier for the customer to automatically retrieve the field values from a customer database.

The **Data Provider Settings** panel of the Scoring Configuration wizard identifies the data provider to use for automatic retrieval of input values. When you use a data provider, the provider must be one of the following objects that are defined in the IBM SPSS Collaboration and Deployment Services Repository:

- A data access plan for an Analytic Data View
- A single Data Provider Definition - Real Time for an Enterprise View.

**Important:** The use of Enterprise View is deprecated. Whenever possible, use an Analytic Data View instead.

After you specify the data provider settings, click **Next** to specify more settings.

## Analytic data view

**Analytic Data View.** Select the data view to use from the list of all available analytic data view definitions that are currently specified within the system.

**Label.** Specify the data view version by selecting the corresponding label from the list of all labels available for the selected analytic data view.

**Data Access Plan.** Select the data access plan to use from the list of plans available in the specified version of the analytic data view.

**Model Input Table.** This list identifies the tables that are used by the scoring model. When you use a data access plan, at least one input table must be associated with a table from the plan. A boldface font for a table name indicates that table is associated with a data access plan table. The values for any input tables that are not associated with a plan table must be supplied manually for each score request.

**Analytic Data View Table.** Defines the table from the data access plan that supplies the values for a selected input table.

Associate an input table with a data access plan table, by performing the following steps:

1. Select the table from the **Model Input Table** list.
2. Select the analytic data view table from the **Table** list.

Repeat these steps for each input table that retrieves data from the data provider.

## Enterprise view

**Data Provider.** Select the definition to use from the list of all available real-time data provider definitions that are currently specified within the system. To be an available data provider, the provider data set must specify a key that matches the key in the table in the Enterprise View. In environments that use many version labels, the system may be unable to retrieve the availableData Provider Definition - Real Time list for selecting a data provider. If this problem occurs, contact the system administrator who should examine the server log for truncation errors. If the number of labels cannot be reduced, the administrator should increase the character limit configuration setting for the repository server. For more information, see the Repository Configuration chapter of the Administrator's Guide.

**Label.** Select the label to use from the list of all labels available in the selected real-time data provider definition.

**Model Input Table.** This list identifies the tables that are used by the scoring model. When you use a data provider, at least one input table must be associated with a table from the data provider. A boldface font for a table name indicates that table is associated with a data provider table. The values for any input tables that are not associated with a data provider table must be supplied manually.

**Data Provider Table.** Defines the table from the data provider that supplies the values for a selected input table.

- **Enabled.** Select this option to associate the selected input table with a data provider table.
- **Table.** Select the data provider table to associate with the selected input table from the list of all tables available in the selected version of the real-time data provider definition.
- **Key.** Select the key to use from the list of all keys available as a result of the real-time data provider definition, label, and table specifications. Keys that are identified as Unique retrieve a single record for each key value. Keys that are identified as Not Unique may retrieve multiple records for each key value.

To associate an input table with a data provider table, perform the following steps:

1. Select the table from the **Model Input Table** list.
2. For the Data Provider Table, click **Enabled**.
3. Select the data provider table from the **Table** list.
4. Select the key to use for retrieving value from the specified table.

Repeat these steps for each input table that retrieves data from the data provider.

## Input data order

The Input Data Order panel of the Scoring Configuration wizard allows control over the order of the input fields for scoring.

The list displays the input fields organized into tables used by the predictive model. Click the plus sign preceding a table name to reveal the input fields for that table. Click the minus sign preceding a table name to hide the input fields for that table.

If a client scoring request omits field names, the input values being passed are assumed to be in the order defined by this panel. Omitting the field names may streamline client and server communications for configurations involving a high volume of scores. In addition, if input fields are included in the scoring response, the order of the returned fields corresponds to the order defined here. See the topic “Input data returned settings” for more information.

To change the order of the input fields within a table, select a field and click the **Up** or **Down** buttons to move the field to the target position. You cannot move a field from one table to another. However, you can reorder the tables by selecting the table name and clicking the **Up** or **Down** buttons.

Click **Next** to specify additional settings.

## Input data returned settings

The Input Data Returned Settings panel of the Scoring Configuration wizard allows control over the contents of the scoring response for inputs.

**Return request inputs in response.** If selected, the scoring response contains values for specified request inputs.

The list displays the input fields organized into tables used by the predictive model. Click the plus sign preceding a table name to reveal the input fields for that table. Click the minus sign preceding a table name to hide the input fields for that table.

The box before each input field and table name indicates whether or not the field or table are included in the response. If the box contains a check mark, the field is included. To add a field to the scoring response, click the empty box before the field name. Click the empty box before a table to add all fields in the table to the response. Clicking a box that contains a check mark removes the corresponding field or table from the response.

Click **Next** to specify additional settings.

## Output data returned settings

The Output Data Returned Settings panel of the Scoring Configuration wizard allows control over the contents of the scoring response for outputs.

Scoring requests can produce a variety of scoring outputs, with the type of output being dependent on the model used for scoring. The scoring configuration allows the specification of a subset of outputs to be returned in response to a scoring request. For models that offer many scoring outputs, restricting the returned values to a small set can help to optimize scoring performance.

The model outputs that can be included in the response appear in the list. For information about specific items for PMML files, see “Supported scoring functions and models” on page 96.

Select the outputs to include for this particular scoring configuration. At least one output item must be selected. Click **Next** to specify additional settings.

## Logging settings

The Logging Settings panel of the Scoring Configuration wizard defines the content of audit logs for the scoring process.

To capture logging information, select **Enable Logging**. Measures that can be included in the logs for scoring appears in the list.

**Request Inputs.** The values of the input fields used for scoring.

**Context Data.** The values of the context fields used for scoring.

**DPD Outputs.** The scoring input values supplied by the Data Provider Definition. These values may be derived from the request inputs and context data.

**Model Outputs.** Output values from the model. The list of available outputs depends on the type of model being configured for scoring. See the topic “Supported scoring functions and models” on page 96 for more information.

### Important:

- For inputs to be logged, logging must be explicitly enabled for each input, including request inputs, context inputs, and DPD outputs. The default is to not log input values.
- Request inputs are logged separately from context inputs. A request input can have the same name as a context input, but they will appear separately in the log.
- If the user does not explicitly provide the input and the data service does not calculate a value for the input, it will be logged as a null value (empty tag).
- The data service will not provide the input unless required by the model. If logging is enabled for an input that is not required, a null value will be logged.

**Scoring Engine Metrics.** Measurements that reflect the performance of the scoring configuration.

Available metrics include:

- **Score Elapsed Time.** Amount of time, measured in milliseconds, between the request for a score and the generation of the score.
- **Score Data Initialization Time.** Amount of time, measured in milliseconds, that the data service takes to be initialized for a score request.
- **Score Data Access Time.** Amount of time, measured in milliseconds, that the data service takes to access its data.
- **Score Computation Wait Time.** Amount of time, measured in milliseconds, that the score provider worker spends waiting for the data service.
- **Score Computation Time.** Amount of time, measured in milliseconds, that the score provider worker spends computing the score.
- **Average Latency.** Average amount of time, measured in milliseconds, between the request for a score and the generation of the score.
- **Minimum Latency.** Smallest amount of time, measured in milliseconds, between the request for a score and the generation of the score.
- **Maximum Latency.** Largest amount of time, measured in milliseconds, between the request for a score and the generation of the score.
- **Score Data Initialization Time.** Amount of time, measured in milliseconds, that the data service takes to be initialized for a score request.
- **Average Data Initialization Time.** Average amount of time, measured in milliseconds, that the data service takes to be initialized for a score request.
- **Minimum Data Initialization Time.** Smallest amount of time, measured in milliseconds, that the data service takes to be initialized for a score request.

- **Maximum Data Initialization Time.** Largest amount of time, measured in milliseconds, that the data service takes to be initialized for a score request.
- **Score Data Access Time.** Amount of time, measured in milliseconds, that the data service takes to access the data.
- **Average Data Access Time.** Average amount of time, measured in milliseconds, that the data service takes to access the data.
- **Minimum Data Access Time.** Smallest amount of time, measured in milliseconds, that the data service takes to access the data.
- **Maximum Data Access Time.** Largest amount of time, measured in milliseconds, that the data service takes to access the data.
- **Score Computation Wait Time.** Amount of time, measured in milliseconds, that the score provider worker spends waiting for the data service.
- **Average Computation Wait Time.** Average amount of time, measured in milliseconds, that the score provider worker spends waiting for the data service.
- **Minimum Computation Wait Time.** Smallest amount of time, measured in milliseconds, that the score provider worker spends waiting for the data service.
- **Maximum Computation Wait Time.** Largest amount of time, measured in milliseconds, that the score provider worker spends waiting for the data service.
- **Score Computation Time.** Amount of time, measured in milliseconds, that the score provider worker spends computing the score.
- **Average Computation Time.** Average amount of time, measured in milliseconds, that the score provider worker spends computing the score.
- **Minimum Computation Time.** Smallest amount of time, measured in milliseconds, that the score provider worker spends computing the score.
- **Maximum Computation Time.** Largest amount of time, measured in milliseconds, that the score provider worker spends computing the score.
- **Average Log Serialization Time.** Average amount of time, measured in milliseconds, to create a log entry in XML format.
- **Minimum Log Serialization Time.** Smallest amount of time, measured in milliseconds, to create a log entry in XML format.
- **Maximum Log Serialization Time.** Largest amount of time, measured in milliseconds, to create a log entry in XML format.
- **Average Log Queue Time.** Average amount of time, measured in milliseconds, to place the XML log data on to the JMS queue.
- **Minimum Log Queue Time.** Smallest amount of time, measured in milliseconds, to place the XML log data on to the JMS queue.
- **Maximum Log Queue Time.** Largest amount of time, measured in milliseconds, to place the XML log data on to the JMS queue.
- **Configuration Scores.** Total number of scores produced by a specific scoring configuration.
- **Score Elapsed Time.** Amount of time, measured in milliseconds, since the previous score generation.
- **Configuration Uptime.** Amount of time, measured in seconds, that the scoring configuration has been available for scoring.
- **Cache Hits.** Number of successful attempts to retrieve data from the memory cache for a scoring configuration.
- **Cache Misses.** Number of failed attempts to retrieve data from the memory cache for a scoring configuration. Each failed attempt results in a new service call to retrieve the necessary data.

**Scoring Engine Properties.** Characteristics of the scoring configuration itself. Entries include:

- **Model Path.** IBM SPSS Collaboration and Deployment Services Repository path for the model file associated with the configuration.

- **Model MIME type.** MIME type of the model file associated with the configuration.
- **Model Version Marker.** Marker identifying the version of the model file associated with the configuration.
- **Model Version Label.** Label identifying the version of the model file associated with the configuration.
- **Scoring Configuration Name.** The name of the scoring configuration.
- **Model ID.** IBM SPSS Collaboration and Deployment Services Repository identifier for the model file associated with the configuration.
- **Scoring Configuration Serial.** Unique identifier of the scoring request for the configuration.

If the configuration uses a data provider, additional scoring engine properties can be included in the scoring logs.

- **Enterprise View Version Marker.** Marker identifying the version of the Enterprise View associated with the data provider.
- **Application View Name.** The name of the Application View associated with the data provider.
- **Application View Path.** IBM SPSS Collaboration and Deployment Services Repository path for the Application View associated with the data provider.
- **Application View ID.** Internal IBM SPSS Collaboration and Deployment Services Repository identifier for the Application View associated with the data provider.
- **Application View Version Marker.** Marker identifying the version of the Application View associated with the data provider.
- **Data Provider Definition Name.** The name of the data provider.
- **Data Provider Definition Path.** IBM SPSS Collaboration and Deployment Services Repository path for the data provider associated with the configuration.
- **Data Provider Definition ID.** Internal IBM SPSS Collaboration and Deployment Services Repository identifier for the data provider.
- **Data Provider Definition Version Label.** Label identifying the version of the data provider.
- **Data Provider Definition Version Marker.** Marker identifying the version of the data provider.

Select the items to include in the log for this particular scoring configuration. For each score request, values for all selected items will be entered into the scoring log.

Click **Next** to specify additional settings.

## Advanced settings

The Advanced Settings panel of the Scoring Configuration wizard offers optional settings used to optimize the scoring process.

**Usable in Batch Scoring.** If selected, the scoring model configuration can be used in batch scoring requests.

**Model Cache Size.** The number of concurrent scoring requests for the configuration. To optimize performance, keep the cache size as low as possible.

**Log Destination.** By default, IBM SPSS Collaboration and Deployment Services uses the Java Message Service (JMS) for logging scoring information to a queue. If your environment is configured to use a custom message driven bean for logging, specify the log destination for that bean. Contact your administrator for the correct destination for your environment.

Click **Finish** to create a scoring configuration with the specified settings.

## Scoring configuration aliases

A scoring configuration alias provides a fixed name to which scoring clients can send scoring requests. The alias passes the requests to an assigned scoring configuration. The scoring configuration that is assigned to the alias can be changed at any time, changing the model that is used for scoring.

Client applications reference the alias name for scoring requests instead of referencing a scoring configuration name. When you assign a different configuration to the alias, subsequent requests are routed to the new configuration. As a result, you can change the configuration that is used for scoring without modifying your client application in any way.

Consider the environment that is illustrated in Figure 7. The IBM SPSS Collaboration and Deployment Services Repository contains two scoring models. *Model 1* is configured for scoring according to the settings that are specified in *Configuration A*. *Model 2* is configured for scoring according to the settings that are specified in *Configuration B*. *Configuration A* is assigned to a configuration alias. The scoring client sends a scoring request to the scoring configuration alias, which sends the request to *Configuration A*. *Model 1* generates the score and sends the result back to the client.

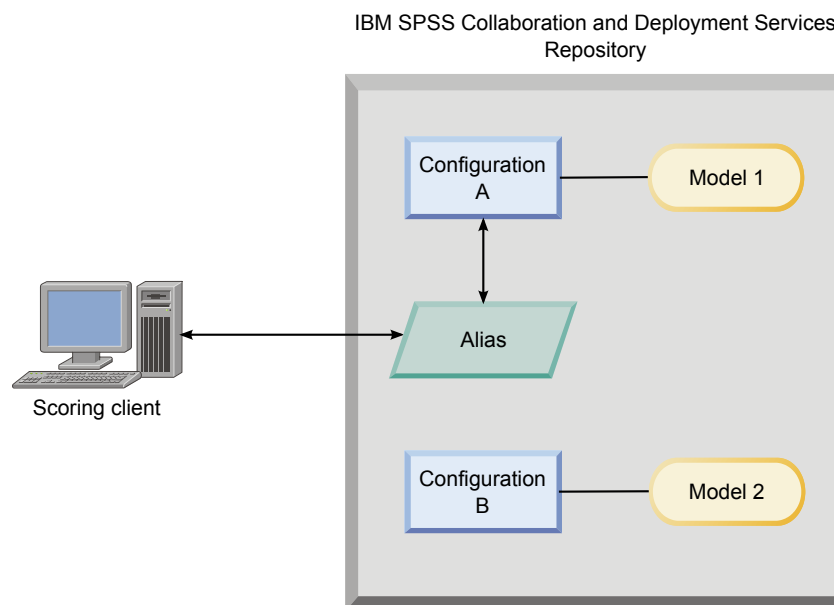


Figure 7. Example alias for scoring configurations

Suppose that you change the configuration that is assigned to the alias to *Configuration B*. In this case, the scoring client still sends the scoring request to the same alias, but the request is then sent to *Configuration B*. *Model 2* generates the score, which is then returned to the client. No change was made to the client, but an entirely different model produced the score. If you did not use the alias, the client would need to be changed to have *Model 2* used for scoring.

Scoring configuration aliases are typically used in A/B testing of scoring models, when you need the ability to replace a production model with a better-performing model without impacting any current production processes.

## Creating scoring configuration aliases

Create an alias for a scoring configuration if you need the ability to change the configuration that is used by scoring clients without modifying the clients.

## Before you begin

Open the Scoring view to access all available scoring configurations and configuration aliases. There must be at least one scoring configuration in the system to create an alias.

## About this task

A configuration alias provides an extra layer between a scoring client and a scoring configuration. You can change the configuration that handles scoring requests without needing to update the client.

## Procedure

1. Select a scoring configuration in the Scoring view.
2. Click **Create Alias**.
3. In the Scoring Configuration Alias dialog box, enter a name for the alias. Scoring tasks reference this name when requests are submitted for scoring.
4. Select the scoring configuration to assign to this alias. This configuration processes all scoring requests that are submitted to the alias.

**Restriction:** An alias cannot reference another alias. Aliases can reference scoring configurations only.

5. Click **OK**.

## Results

The Scoring view updates to include an entry for the alias. The configuration alias is available for scoring requests. The alias passes all submitted scoring requests to the specified scoring configuration.

## Editing scoring configuration aliases

Change the scoring configuration that is referenced by an alias to use a different configuration for generating scores for requests that are submitted to the alias.

## Before you begin

Open the Scoring view to access all available scoring configurations and configuration aliases.

## About this task

When you use a configuration alias to route scoring requests, you can change the scoring configuration that is used to generate scores by modifying the alias. This approach is typically done when a new scoring model outperforms another model. Update the alias to reference a scoring configuration for the new model to have all subsequent scoring requests that are submitted to the alias passed to the new scoring configuration.

When you assign a new configuration to an alias, the model that is used by the new configuration must be compatible with the model that is used by the old configuration. The inputs for the new model must be a subset of the inputs for the old model. The new model outputs must be a superset of the old model outputs. In addition, the input field types must match between the two models. For example, if the old model uses the fields *age*, *gender*, and *income* to produce predicted target values, probabilities, and confidence values, the new model must use at least one of the same input fields to produce output that contains these three prediction values at a minimum. If *gender* has a type of string in the old model, it must also be a string type in the new model. These compatibility restrictions ensure that the input from the client score request can be processed by both models and the results can be handled properly by the client. If the models are not compatible, changes to the scoring client are required when a new configuration is assigned to an alias.



## Procedure

1. Select a scoring configuration alias in the Scoring view.
2. Click **Edit Alias**.
3. Select a new scoring configuration to assign to the alias.
4. Click **OK**.

## Results

The Scoring view updates to show the new configuration that is assigned to the alias. The modified alias is available for scoring requests.

## Deleting scoring configuration aliases

Delete a configuration alias if it is no longer needed for scoring tasks.

### Before you begin

Open the Scoring view to access all available scoring configurations and configuration aliases.

### About this task

If an alias is no longer needed, delete the alias from the system. The scoring configuration that is used by an alias is not deleted when the alias is deleted. If the configuration is no longer needed, delete the configuration separately.

**Important:** Verify that no client applications are currently using a configuration alias before you delete the alias. If you delete an alias that is used by a client application, that application can no longer generate scores.

## Procedure

1. Select the alias in the Scoring view.
2. Click **Delete Configuration** or press the Delete key.

## Results

The alias is removed from the system.

---

## Scoring configuration association sets

A scoring configuration association set consists of a primary scoring configuration and one or more alternative configurations for generating scores. Scoring requests that are submitted to the primary configuration are assigned to one of the configurations in the set according to specified distribution percentages.

Use a scoring configuration association set to compare the performance of one or more models with a current production model. Scoring requests that are submitted to the production model are automatically assigned to one of the alternative models for processing. Over time, you can compare the scores that are generated by all of the models in the set to determine which one performs best.

You can stop testing of all alternative models by deleting the association set. When the set is deleted, the primary configuration handles all scoring requests. Alternatively, you can stop testing individual alternative models by removing them from the set and adjusting the distribution percentages.

Consider the environment that is illustrated in Figure 8 on page 108. The IBM SPSS Collaboration and Deployment Services Repository contains an association set that includes three scoring configurations.

*Configuration B* is the primary configuration and contains the scoring settings for *Model 2*. *Configuration A* and *Configuration C* define the scoring settings for *Model 1* and *Model 3*. The association set sends 15% of score requests to *Configuration A* and 15% of score requests to *Configuration C*. The remaining 70% of score requests are processed by *Configuration B*.

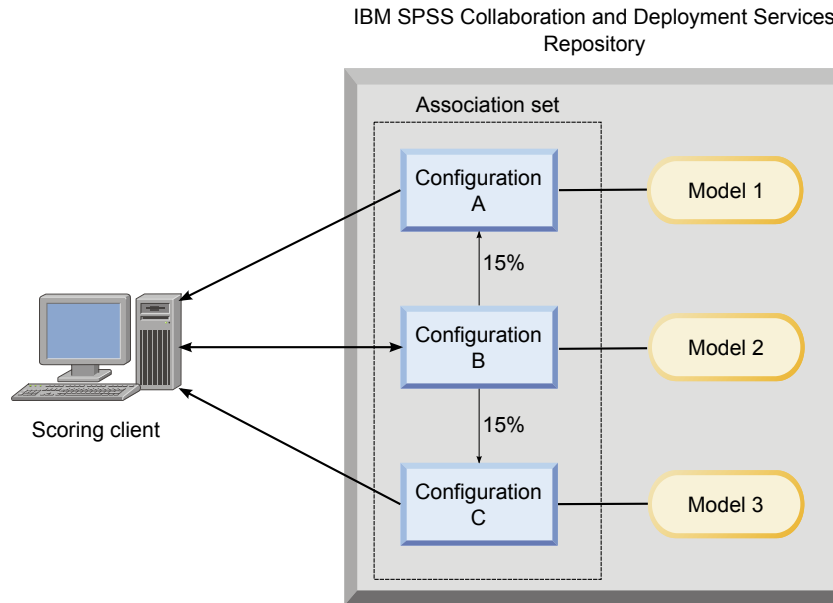


Figure 8. Example association set for scoring configurations

When a scoring client sends a scoring request to *Configuration B*, it is routed to one of the configurations in the set. That configuration generates the score and returns the result to the client.

If the configuration that is assigned to a scoring request is deleted from the system, is not running, or is experiencing an error that makes it unavailable, the request is rerouted to the primary configuration for the association set.

## Configuration compatibility

To be included in an association set, a scoring configuration must be compatible with the primary configuration for the set. Compatibility ensures that score requests sent to the primary configuration can be processed by any associated configurations. In addition, the outputs that are returned by associated configurations can be processed by the client application that sends the score request.

To be compatible, the following criteria must be satisfied:

- All input table names in an associated configuration must have corresponding table names in the primary configuration
- The inputs for an associated scoring configuration must be a subset of the primary configuration inputs
- The outputs for an associated scoring configuration must be a superset of the primary configuration outputs

For the associated scoring configuration inputs to be a subset of the primary configuration inputs, the following criteria must be satisfied:

- Names of fields that are required in the input tables for an associated configuration must be included in the corresponding tables of the primary configuration. The fields must be required in the primary configuration.

- Names of fields that are optional in the input tables for an associated configuration do not need to be included in the tables of the primary configuration. However, if those fields are included in the primary configuration, they can be either required or optional.
- The types for the input fields in the associated configuration must match the types for the input fields in the primary configuration

Figure 9 illustrates both compatible and incompatible scoring configurations for a primary configuration.

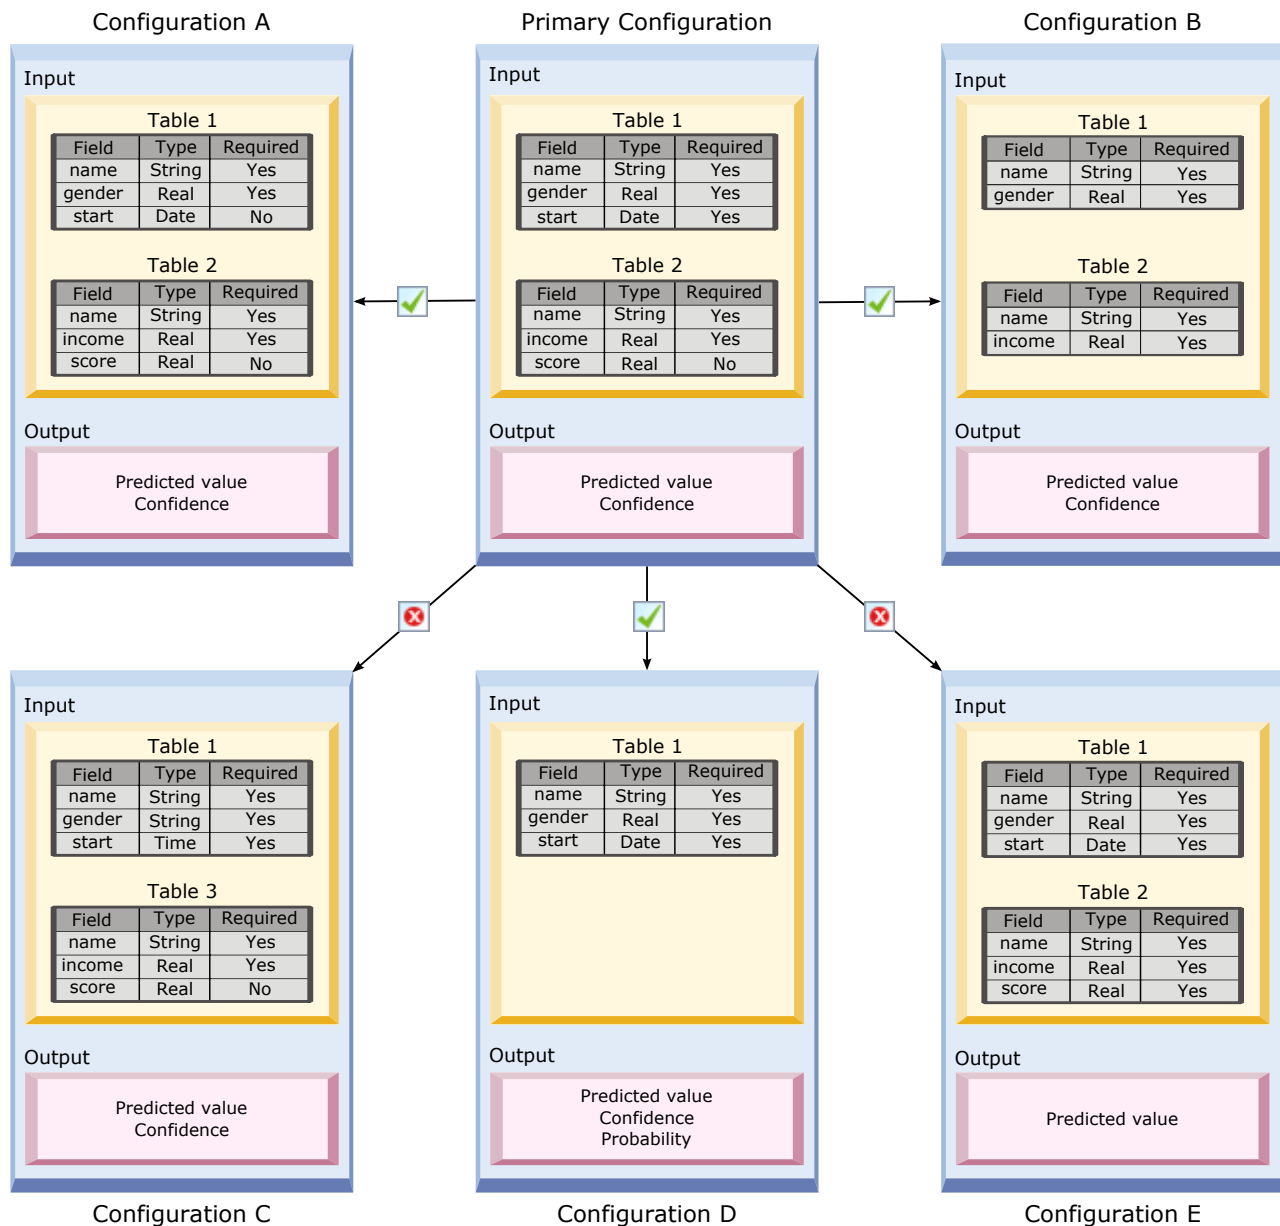


Figure 9. Scoring configuration compatibility

Configuration A and Configuration B are both compatible with the primary configuration. These configurations have two tables that have the same names as the tables in the primary configuration. All of the fields in the configurations have the same names and types as the fields in the primary configuration. All of the fields that are required for the associated configurations are also required in the primary configuration. Finally, all of the configurations return the same set of outputs.

*Configuration D* is also compatible with the primary configuration. The one input table in *Configuration D* has a corresponding table in the primary configuration. All of the fields in the associated configuration have the same names and types as the fields in *Table 1* of the primary configuration. The outputs of *Configuration D* are also a superset of the outputs for the primary configuration.

In contrast, *Configuration C* is not compatible with the primary configuration for two reasons. First, the primary configuration does not include a table that is named *Table 3*. A score request that is sent to the primary configuration would not include any data for this table. Second, the *start* field is of the Time type for *Configuration C* but has a type of Date in the primary configuration. The data for *start* is of different types in the two configurations.

*Configuration E* is also not compatible with the primary configuration for two reasons. First, the *score* field is required for the associated configuration but it is optional for the primary configuration. So, a score request that is sent to the primary configuration might not include data that is required for the associated configuration. Second, *Configuration E* returns only the predicted value while the primary configuration returns both the predicted value and the confidence. The associated configuration outputs are not a superset of the primary configuration outputs.

## Creating scoring configuration association sets

Create a scoring configuration association set to distribute score requests across a set of configurations.

### Before you begin

Open the Scoring view to access all available scoring configurations and configuration aliases. There must be at least two scoring configurations in the system to create an association set.

### Procedure

1. Select a scoring configuration in the Scoring view. This configuration is the primary configuration for the association set.
2. Click **Create Associations**.
3. In the Scoring Configuration Association dialog box, select one or more scoring configurations to associate with the primary configuration. The list includes only the configurations that are compatible with the primary configuration. Press the **Ctrl** key to select multiple configurations. To select a contiguous set of configurations, press the **Shift** key and select the first and last configurations in the set.
4. Click **Next**.
5. Enter the percent of scoring requests to submit to each configuration by typing the value in the **Percentage** column. The specified percent values must sum to 100%.
6. Click **Finish**.

### Results

The Scoring view updates to show the new configuration association set. Scoring requests sent to the primary scoring configuration in the association set are distributed across the configurations in the set in accordance with the specified percentages.

## Editing scoring configuration association sets

Modify a scoring configuration association set to change the configurations included in the set or to change the percentage of requests that are sent to each configuration in the set.

### Before you begin

Open the Scoring view to access all available scoring configurations and configuration aliases.

## Procedure

1. In the Scoring view, select the primary configuration for an association set.
2. Click **Edit Associations**.
3. In the Scoring Configuration Association dialog box, select one or more scoring configurations to associate with the primary configuration. The list includes only the configurations that are compatible with the primary configuration. The configurations currently included in the association set are selected by default. Press the **Ctrl** key to select multiple configurations. To select a contiguous set of configurations, press the **Shift** key and select the first and last configurations in the set.
4. Click **Next**.
5. Enter the percent of scoring requests to submit to each configuration by typing the value in the **Percentage** column. The specified percent values must sum to 100%.
6. Click **Finish**.

## Results

The Scoring view updates to show the modified configuration association set. Scoring requests sent to the primary scoring configuration in the association set are distributed across the configurations in the set in accordance with the specified percentages.

## Deleting scoring configuration association sets

Remove the associations from a scoring configuration to prevent scoring requests that are submitted to the configuration from being processed by any associated configurations.

### Before you begin

Open the Scoring view to access all available scoring configurations and configuration aliases.

## Procedure

1. In the Scoring view, select the primary scoring configuration in an association set whose associations you want to remove.
2. Click **Remove Associations**.

## Results

The associations between the selected scoring configuration and other configurations are removed from the system. All subsequent scoring requests that are submitted to the selected configuration are processed by that configuration.

---

## A/B split testing for scoring configurations

Use A/B split testing to compare different scoring configurations in terms of overall performance. You can add the best configuration to your existing production tasks to improve their predictive capabilities.

When you want to use a model to generate scores, you define a scoring configuration that is based on a particular version of the model in the IBM SPSS Collaboration and Deployment Services Repository. Over time, you might need to update the model to take into account extra considerations, such as new data availability or better predictive algorithms. However, before you replace a model that is used in production tasks, you might want to evaluate the performance of the new model to determine whether it is better than the current model.

In A/B split testing, you keep the existing scoring configuration deployed but reroute some of the incoming score requests from the production configuration to the potential replacement configuration. As more requests are processed, you can review the model results to evaluate the overall performance. If you

determine that the replacement configuration is working better than the existing configuration, replace instances of the existing configuration in your production tasks with the new configuration.

The A/B split testing approach incorporates both association sets and configuration aliases. The association set defines the scoring configurations that are being evaluated and the number of requests that are sent to each. The alias controls the configuration that is used in your production tasks.

To illustrate this approach, suppose that you generate scores from Model 2 by using Configuration B. You develop two possible alternative models for generating scores, Model 1 and Model 3, and create a scoring configuration for each. You want to determine whether either alternative model performs better than Model 2.

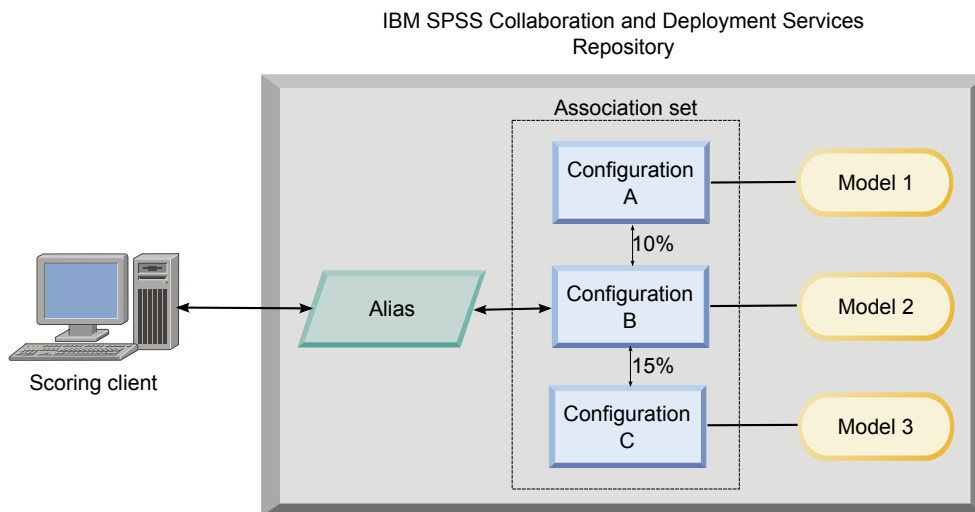


Figure 10. A/B testing example

Figure 10 displays an A/B split testing environment for evaluating the configurations. A client application sends a score request to a configuration alias, which passes the request to Configuration B. Configuration B is the primary configuration in an association set that sends 10% of requests to Configuration A and 15% of requests to Configuration C. The remaining 75% of requests are processed by Configuration B. The configuration that is assigned to the request generates and returns the score output by using the associated model.

The scoring log table contains the results for all of the scoring configurations. After you examine the results, suppose that you determine that Configuration C produces better overall scores than Configuration B. To deploy Configuration C in your production tasks, edit the alias to reference this configuration. All subsequent requests that are sent to the alias are routed to the new configuration.

If none of the alternative configurations produce better results, remove the association set to have the existing configuration continue to produce scores in your production tasks.

## Evaluating scoring models by using A/B testing

For A/B split testing, use a scoring configuration association set to define the set of configurations that are included in the tests. Use a configuration alias for the primary configuration in the set to manage the configuration that is used for production-level scoring tasks.

### Procedure

1. Create an alias for the primary scoring configuration. Client applications send score requests to this alias.

2. Create an association set for the primary scoring configuration that includes the other configurations to be tested. Assign the percentages of score requests to be handled by each configuration. Scoring requests that are passed to the association set by the alias are distributed across the configurations in the set.
3. Send score requests to the alias and review the results. Use the information in the scoring log table to evaluate the performance of the configurations that are being tested.
4. If the results indicate that a scoring configuration produces better results than the primary configuration, update the alias to reference the new configuration. Client applications that send requests to the alias automatically use the new configuration to generate scores.
5. Delete the association set that was used for testing.

---

## Scoring view

The Scoring view lists all of the scoring configurations, configuration aliases, and configuration association sets currently defined in the system. From this view, aliases and associations can be created. In addition, configurations, aliases, and associations can be modified or deleted from the system.

To access the Scoring view from the menus, select:

**View > Show View > Scoring**

By default, configurations on the current server appear in the view. To access other servers, select the server from the **Server** list. If you are not logged in to the selected server, the Login dialog box opens, prompting you to log in to the server.

Each row in the view corresponds to a scoring configuration or configuration alias for the server. Properties that are shown for each entry include the following values:

**Type** An indicator of the type of model that is used by the configuration.

**Status** The status for the element.

**Configuration Name**

The name of the scoring configuration or alias. For an alias, the assigned scoring configuration is displayed in parentheses after the alias name.

**Model** The name of the model file that is associated with the configuration.

**Label** The label that identifies the version of the model file that is configured for scoring.

**Average Score Time**

The average time for a single score to be calculated.

**Scores/Second**

The number of scores that are produced per second at the time the scoring value is retrieved.

**Total Scores**

The total number of scores that are generated by using the configuration.

If a configuration is the primary configuration in an association set, the row begins with a triangle icon that indicates the entry can be expanded or collapsed. When you click the triangle, the row expands to show all of the configurations included in the set. The percentage of scoring requests that are assigned to each configuration is displayed in parentheses after the name. If you click the triangle for an expanded row, the row collapses to hide the association set details.

The configuration list can be refreshed to update the scoring statistics. To refresh the view, click the **Refresh** icon.

The list can be sorted by any column in the view. Click a column heading to sort by that column. Click the heading again to reverse the sort direction.

## Filtering the Scoring view

The Scoring view can be filtered to reduce the number of results that appear in the list.

Once filtering is enabled, the specified filters persist throughout the IBM SPSS Collaboration and Deployment Services Deployment Manager session until changed. In addition, filter settings persist across server connections. For example, if filtering is enabled and the server selection is changed from server A to server B, the filter settings established for server A persist for server B. To access view filters, click the Filters button. The Scoring Filters dialog box opens.

**Enable Filtering.** If selected, the specified filters are applied to the view.

The view can be filtered by the configuration name. Select the configurations to view in the Scoring view.

## Editing a scoring configuration

To edit a scoring configuration:

1. In the Scoring View, right-click the configuration to be modified.
2. Select **Edit** from the pop-up menu. The Scoring Configuration wizard appears.
3. Modify the configuration settings as needed.
4. Click **Finish**.

## Suspending and resuming scoring configurations

At times, it may be necessary to temporarily suspend a scoring configuration from processing requests. Suspending a configuration frees associated system resources and allows various maintenance tasks to be performed without affecting incoming score requests. Resuming a suspended configuration allows the configuration to begin processing scoring requests again.

To suspend a scoring configuration:

1. In the Scoring View, select the configuration to be suspended. To select multiple configurations, press the Ctrl key while selecting.
2. Right-click a selected configuration and select **Suspend Configuration(s)** from the pop-up menu. Alternatively, press the **Suspend Configuration(s)** button on the scoring view.

The status for the selected configurations updates accordingly.

To enable a suspended configuration to accept scoring requests, the configuration must be resumed. To resume a scoring configuration:

1. In the Scoring View, select the configuration to be resumed. To select multiple configurations, press the Ctrl key while selecting.
2. Right-click a selected configuration and select **Resume Configuration(s)** from the pop-up menu. Alternatively, press the **Resume Configuration(s)** button on the scoring view.

The status for the selected configurations updates, indicating the configurations are active.

## Deleting a scoring configuration

To delete a scoring configuration:

1. In the Scoring View, select the configuration to be deleted. To select multiple items, hold down the Ctrl key and select additional rows.
2. Click the **Delete** button. The Delete Confirmation dialog box appears.



3. Click **OK**. The configuration is removed from the system.

Alternatively, right-click the configuration and select **Delete** from the pop-up menu.

Deleting a scoring configuration does not remove the model associated with the configuration from the system. Furthermore, any other configurations using the model remain in the system until specifically deleted.

---

## Scoring Graph view

The Scoring Graph view displays the scoring throughput for a selected scoring model.

To access this view:

1. In the Scoring View, select the configuration to be viewed as a graph.
2. Click the **Graphic View** button.

Alternatively, right-click the configuration and select **Graphic View** from the pop-up menu. The Scoring Graph view for the configuration opens.

The appearance of the graph is defined by the following controls:

**Content Shown.** Select the scoring metric appearing in the graph. Available metrics include the following:

- **Service Scores.** Total number of scores produced by all configurations.
- **Service Uptime.** Amount of time, measured in seconds, that the scoring service has been available for scoring.

For a description of other available measurements, see the Scoring Engine Metrics section of “Logging settings” on page 102.

**Show Latest Period (Hour).** Limits the graph to the most recent time period. Select the size of the period as the number of hours to show. For example, to show the past half-hour, specify a value of *0.5*. To show the past two hours, specify a value of *2*.

The graph in the view is regenerated at regular intervals, with the most recent information taking precedence over older results.

The Scoring Graph view displays the performance for a single selected configuration. To monitor multiple configurations, open multiple graph views.



---

## Chapter 10. Jobs

---

### What is a job?

A job is a container for a set of steps. Each step has parameters associated with it. Before you execute a step, you must embed it within a job. In order to generate results, a job must contain at least one step. (Although an empty job will run, it does not produce any results.) Jobs are scheduled and executed by an application server.

Job steps can be executed sequentially or conditionally. For example, you can make the second step in a job contingent upon the outcome of the first step in a job. Objects stored in the repository can provide input for a job step, and job results can be stored in the repository. For example, a job might include a data preparation step, which reads data from a data mart, followed by a step in which an IBM SPSS Modeler stream reads the prepared data and computes a propensity score based on the data.

Jobs are created in the Content Explorer and modified in the Job Editor. Jobs are stored in the database associated with the repository.

You can move jobs anywhere within the Content Repository. However, unlike other files, jobs are native and unique to the system. You must create and work with the job in the IBM SPSS Collaboration and Deployment Services Deployment Manager.

Jobs are executed on a server. The content of the job determines the type of execution server required. For example, if a job contains an IBM SPSS Modeler stream, an IBM SPSS Modeler server is required. The server can reside on the same system as the IBM SPSS Collaboration and Deployment Services server. Alternatively, certain jobs can be executed on remote servers.

### Versioning and labeling jobs

Just like any other repository object, jobs can be versioned. In addition, a job can have many schedules associated with it. However, although a job can have many schedules, a schedule can contain only one job.

The following guidelines apply:

- Each time that changes are saved to a job, a new version of the job is created.
- Only a labeled version of a job can be scheduled.
- The system applies the *LATEST* label to the most recent version of a job. If a label is not explicitly supplied, the only option when scheduling a job is the version labeled *LATEST*.

### Components of a job

A job can contain any combination of the following components:

- Report files. See the topic Chapter 14, “Reporting job steps,” on page 157 for more information.
- Visualization report files. See the topic Chapter 15, “Visualization report job steps,” on page 163 for more information.
- SAS syntax files. See the topic Chapter 16, “SAS® job steps,” on page 167 for more information.
- General job steps. See the topic Chapter 17, “General job steps,” on page 173 for more information.
- Message-based job steps. See the topic Chapter 18, “Message-based job steps,” on page 183 for more information.
- Notification job steps. See the topic Chapter 19, “Notification job steps,” on page 185 for more information.

In addition, products that collaborate with IBM SPSS Collaboration and Deployment Services may offer additional components that can be included in a job, including:

- IBM SPSS Modeler streams.
- Scenario files.
- Champion challenger job steps.
- IBM SPSS Statistics syntax files.

To include these components, install the necessary adapters and plug-ins as described in the documentation for the product offering the additional functionality.

## Prerequisite to running jobs

For certain types of jobs, you must establish server and credential definitions.

- Server definitions are required for IBM SPSS Modeler streams, IBM SPSS Statistics syntax, and SAS syntax job steps.
- Credential definitions are required for IBM SPSS Modeler stream job steps.

See the topic Chapter 6, “Resource definitions,” on page 47 for more information.

## External file dependencies

Job steps often refer to external sources, such as data files or databases. For those jobs to complete successfully, the external sources must be accessible to the IBM SPSS Collaboration and Deployment Services Repository.

For example, consider the following IBM SPSS Statistics syntax that references the data file *Employee data.sav*:

```
GET FILE='C:\Program Files\data\Employee data.sav'.  
GRAPH  
/BAR(SIMPLE)=MEAN(salary) BY jobcat .
```

If this syntax is used as an IBM SPSS Statistics job step, the data file must exist in the *C:\Program Files\data* directory on the IBM SPSS Collaboration and Deployment Services Repository computer. If not, the job will fail.

An alternative approach for external files involves storing them on a network location that the IBM SPSS Collaboration and Deployment Services Repository can access. This could involve mapping to a shared drive on a remote machine or using a UNC file reference.

Another method for dealing with external files involves storing the files in the repository. Use a General job step to extract them to a location necessary for subsequent steps. For example, if the repository contains *Employee data.sav*, a General job step can place it in the *C:\Program Files\data* directory for a subsequent IBM SPSS Statistics step.

Finally, if steps make use of databases, an ODBC source for the database must be defined on the IBM SPSS Collaboration and Deployment Services Repository computer.

## Job process overview

Although the components of a job vary, the basic process for working with jobs consists of the following tasks:

1. Create a new job or open an existing job.
2. Add steps to the job.
3. Specify the relationship between job steps, if applicable.
4. Save the job.

5. Test the job by running it immediately (optional).
6. Schedule the job for execution or run the job immediately.
7. Specify email notifications (optional).
8. Save changes to the job.
9. View job status (optional).
10. View output results from the job (optional).

---

## Working with jobs in the Content Explorer

Typically, the Content Explorer is the starting point and the end point for jobs in the system. When you first work with a job, you either create or select the job in the Content Explorer. After you have created and saved the job, you can access it again in the Content Explorer.

### Creating new jobs

Before you can schedule any step for execution, you must create a job. You can create an empty job and add steps to it later; however, to execute a job, your job must contain at least one step.

To make jobs easier to find, it is recommended that you have a designated place for storing them. For example, you can create a separate folder for jobs, or you can store jobs in the same folder as their respective model files. You must save a job within a folder. You cannot save a job within another object.

As you create jobs, it is important to note the following:

- By default, when you create a job, it is automatically saved in the repository. You need to explicitly save the job again after you add steps to it.
- If the Job option is disabled when you try to create a job, this means that you selected an item in the Content Explorer that is not a folder. For example, if you click a IBM SPSS Modeler stream file, the **New > Job** option is disabled because you cannot store a job within a IBM SPSS Modeler stream.
- Although you can house jobs almost anywhere within the Content Explorer, you cannot store jobs within the Resource Definitions folder. See the topic Chapter 6, “Resource definitions,” on page 47 for more information.

To create a new job:

1. In the Content Explorer, create a folder to house your job, if it does not already exist.
2. Select the folder in which you want to store your job.
3. From the File menu, choose:  
**New > Job**  
The New Job Information dialog box opens.
4. In the New Job Information dialog box, provide the following information:

Table 10. New job information.

Field	Description	Required or optional
Name	You must specify a unique name for your job. The name must be unique within the folder. See the topic “Naming conventions within the system” on page 10 for more information.	Required
Description	This field contains a description of your job. Although this field is optional, it is recommended that you provide a detailed description for your job because the description can be searched. For example, if your description contains the phrase <i>propensity analysis</i> , and you search for that term, this job appears in the search results.	Optional

Table 10. New job information (continued).

Field	Description	Required or optional
Keywords	The keywords that you specify here are used by the search facility. See the topic “Searching” on page 17 for more information.	Optional
Expiration date	An expiration date can be specified for the job. See the topic “Working with Expiration Dates and Expired Files” on page 31 for more information.	Optional

5. After you have specified the job details, click **Finish**. The new job appears in the folder that you selected previously in the Content Explorer and as a tab in the Job Editor.

You are now ready to add steps to your job. See the topic “Adding steps to a job” on page 122 for more information.

## Opening an existing job

To open an existing job:

1. In the Content Explorer, navigate to the job and double-click on it. The steps within the job appear in the Job Editor.

## Viewing job properties

Properties for jobs are similar to properties for other objects in the Content Explorer. See the topic “Viewing object properties” on page 25 for more information.

---

## Working with the job editor

The Job Editor is where you modify jobs. You can perform the following tasks:

- Add steps to a job.
- Specify parameters for a job step.
- Establish relationships between steps.
- Schedule a job for execution.
- Specify notifications.
- View job status.

The Job Editor is divided into the following sections:

**Job canvas.** The job canvas provides a visual representation of your job. You add steps to your job in the job canvas. You can drag files from the Content Explorer and select job step types from the Job Steps list. In addition, you can establish relationships between items within a job.

**Job palette.** The job palette is divided into two main sections:

- **Relationships.** Using the tools in the Relationships section, you can establish relationships between the steps in your job. See the topic “Specifying relationships in a job” on page 123 for more information.
- **Job Steps.** In the Job Steps section, you can add several types of job steps to your job—for example, champion challenger, general, message-based, or notification job steps.

**Job properties section.** This section is divided into the following tabs:

- **General information.** Describes general attributes associated with the job. See the topic “General information tab” on page 121 for more information.

- **Job Variables.** Defines variables for the job whose values can be passed to steps within the job. See the topic “Job variables” for more information.
- **Notifications.** Allows you to specify email notifications for job failure and job success. See the topic “Job success and failure notifications” on page 147 for more information.

The General information, Job Variables, and Notifications tabs apply to the entire job. Each step within a job has its own tabs associated with it. When you click a job step, the job processing tabs change accordingly.

## General information tab

When you first open a job, the General Information tab appears by default. You cannot modify any items within the General Information tab. The information on this tab is generated and updated by the application, based on information that is provided at other points in the job creation process.

The General Information tab contains the following details about the job:

**Job Path.** This field specifies the directory location for the job.

**Job Label.** This field contains the label that has been applied to the job. If no label has been applied, this field remains blank.

**Schedule Status.** This field describes whether or not the job has already been scheduled. Valid values are *Scheduled* or *Not scheduled*.

**Last Run Status.** This field describes whether the job ran successfully or failed the last time it was run.

## Job variables

Job variables define parameters whose values can be passed to any step within the job. Using variables, any job can be used as an iterative consumer, in which values external to the job can be used to control job processing. Values for the variables can be defined:

- When initiating the job
- In schedules associated with the job
- In other jobs executing before the job

The Job Variables tab for a job displays a table identifying variables defined for a job.

**Variable Name.** Lists the names of variables defined for the job.

**Default Value.** Identifies the default value for each job variable. If a variable has no specified default value and a value has not been assigned in some other fashion, the user is prompted for a value during job execution.

**Description.** Lists informative text about each variable typically used to aid in identifying the variables.

Job variables can be used in any job step field that supports entry field content assist. In addition to system properties, the list of available variables appearing when typing a \$ in those fields includes all variables defined at the job level.

## Adding variables to a job

The Add Job Variable dialog box creates new job variables. To define a new variable for a job, click **Add** on the Job Variables tab for an open job.

1. Type a unique name for the variable. No two variables for the same job can have the same name. Job variable name can contain English letters and digits only, and the first character must be a letter.

2. Type the value of the variable to use as the default.
3. Type a description for the variable.
4. Click **OK**.

The new variable appears in the variable list for the job.

## Editing existing job variables

To modify an existing variable for a job:

1. Open the job.
2. Select the Job Variables tab.
3. In the variables table, select the cell containing the value to be changed. The variable name cannot be modified.
4. Modify the value.
5. Press the Enter key.

The updated information appears for the variable in the list.

## Removing variables from a job

To remove an existing variable from a job:

1. Open the job.
2. Select the Job Variables tab.
3. Select the variable to remove.
4. Click **Remove**.

The variable is removed from the job and no longer appears in the variable list.

---

## Adding steps to a job

After you have created or selected a job, you are ready to add steps to it. See the topic “Components of a job” on page 117 for more information.

To add steps to a job:

1. Open the job. The job appears in the Job Editor.
2. Select the item that you want to add to the job. You have the following options:
  - **Add a file.** From the Content Explorer, click the file you want to add, and then drag the file into the job canvas. Alternatively, you can right-click on the file and select **Add to Job**.
  - **Add a job step.** From the Job Steps section of the job palette, select the type of job step that you want to add. The option you select will be highlighted. Then click anywhere in the job canvas. By default, the system numbers steps in the order in which they appear. Thus, subsequent job steps appear as Event 2, Event 3, and so on. You can rename a step at any time. The number has no effect upon the order in which steps are executed.
  - **Copy and paste an existing job step.** Right-click an existing job step from any job and select **Copy**. Right-click the canvas for the target job and select **Paste**. Note that job step notifications are not preserved when copying steps.
3. If you want to add additional steps, return to step 2.

If you have two or more steps in your job, you can either:

- Run job steps concurrently. For more information, see “Executing job steps concurrently” on page 123.
- Establish relationships between the steps. For more information, see “Specifying relationships in a job” on page 123.



Many job steps reference IBM SPSS Collaboration and Deployment Services Repository resources that may have multiple versions. These steps reference a particular version of the resource by label. You must have adequate authority to the label selected for the job step to execute the job. If your authority does not permit referencing the job step resource by the selected label when the job runs, the system will be unable to find the resource and the execution will fail. See the topic “Label security” on page 34 for more information.

---

## Executing job steps concurrently

If you create a job with multiple steps but you do not specify a relationship between the steps, the system will execute the steps concurrently. Thus, the steps are not contingent upon one another.

---

## Specifying relationships in a job

When you add steps to a job, you have the option of establishing a relationship between the steps. The relationships that you specify determine the sequence of steps in the job and the conditions under which each step is executed.

Relationships are established by relationship connectors. When you connect two steps, there is an antecedent step and a consequent step. The antecedent step is the starting point for the relationship. This step, as its name suggests, is executed first. The consequent step is contingent upon the completion of the antecedent step. It is important to note that completion does not always mean successful completion. For example, if you specify a fail relationship between two steps, the consequent step is executed only if the antecedent step fails to complete successfully.

For example, suppose you have four steps in your job: A, B, C, and D. A simple scenario for the four steps would be to execute them sequentially--that is, execute A first, then B, then C, and then D. In another case, you might want to execute step A and then execute step B only if step A succeeds. If step A fails, you might want to execute step C. In both cases, you would want to generate a report in step D.

If you cancel a job, the relationships that you establish will not be followed. For example, suppose that you have a job with steps A and B with a sequential relationship between them. If you run the job and cancel it while step A is running, step B will not execute. This behavior holds true for all relationship types, including a fail relationship.

When you create a job, you can establish the following relationships between steps in the job:

*Table 11. Relationship connectors.*

Relationship connector	Conditional expressions allowed?
Sequential	No
Pass	No
Fail	No
Conditional	Yes. In order to use a conditional connector, you must specify at least one argument.

Although the relationships that you can establish between steps vary, the process for establishing a relationship is the same.

To specify a relationship between two steps in a job:

1. Open an existing job. The job appears in the Job Editor.
2. From the palette, select the relationship you want to establish. The relationship is highlighted on the palette.

3. In the job canvas, click the antecedent step in the sequence and then click the consequent step in the sequence. When the mouse is released, a directional arrow connects the two steps. The icon in the center of the arrow describes the relationship.
4. To establish another relationship within the job, return to step 2.
5. Save your changes.

## Sequential connector

If you connect two steps using the sequential connector, the consequent step is executed immediately after the antecedent step is completed.

The sequential connector is always executed. Unlike pass, fail, or conditional relationships, there are no conditions associated with a sequential connector. The only criterion for the execution of the consequent step is the completion of the antecedent step.

## Pass connector

The pass connector is executed if the antecedent step is executed successfully. If the antecedent step is completed successfully, the system proceeds to the consequent step. If the antecedent step fails to complete successfully, the system will not execute the consequent step.

## Fail connector

The fail connector is executed if the antecedent step does not execute successfully. If the antecedent step is executed successfully, the consequent step in a fail relationship is not executed. For example, the consequent step in a fail relationship may generate a report indicating that the antecedent step failed.

## Conditional connector

The conditional connector is executed if the antecedent step satisfies the condition(s) you specify. For conditional relationships, you must specify an expression that contains the condition. This is the only relationship connector in which you can pass arguments to the system.

The Conditional Connector Expression can be any script the system can evaluate, but there are two special parameters that relate to the consequent step:

**Completion\_code.** The `completion_code` is an integer value. When the `completion_code` expression is used, it must be lowercase. The job step interprets the meaning of the completion code. For the General job step, the completion code is the executable return code. For example, for most Windows commands, `File Not Found = 2`.

**Success.** The success parameter is a Boolean, where `true` indicates that the job step execution was successful. When the success expression is used, it must be lowercase.

You can compare and combine parameters, using standard scripting operators. For example,

- `&&`
- `||`
- `==`
- `!=`
- `<`
- `>`
- `!`
- `()`

The following are examples of expressions for a conditional connector:

```
completion_code == 0 && success == true  
completion_code == 1 || success == false  
completion_code >= 10 || success == false  
(completion_code >= 1 && completion_code <= 5) && success == true
```

## Warning expressions in conditional connectors

A special case of using conditional job step connectors is specifying a relationship based on the warning code returned by the antecedent step.

The following job step types support warning codes:

- General
- IBM SPSS Modeler
- Scenario
- IBM SPSS Statistics
- SAS

To specify a warning-code-based relationship:

1. Connect the steps with a conditional connector.
2. Enter `warning==true` as the conditional expression. When the warning expression is used, it must be lowercase.
3. Open the parent job step and override the default warning code on the General tab if necessary.

Once the antecedent step is completed, the system will evaluate the conditional expression and proceed to the next step if `true` is returned.

## Viewing general relationship properties and modifying relationships

Each object in the job canvas has properties associated with it. The relationship between steps is described on the General tab. At times, you may want to modify the relationship between steps in a job.

**Viewing relationship properties.** To view the properties for a relationship:

1. In the Job Editor, click the relationship icon. Two small black boxes appear at each end of the arrow to indicate that the connector has been selected. The General tab appears in the job properties section.

**Modifying relationship properties on the General tab.** To modify a relationship:

1. Select a new relationship from the Relationship drop-down list. The connector icon changes to reflect the new relationship.
2. If the new relationship requires that you provide parameters, you must specify the parameters. For example, if you change a pass connector to a conditional one, you must specify the parameters that describe the condition. See the topic “Conditional connector” on page 124 for more information.
3. Save your changes.

**Modifying relationship properties in the job canvas.** Alternatively, you can modify relationships within the job canvas. To do so:

1. In the Job Editor, right-click the arrow connecting the two steps whose relationship you want to modify. Two small black boxes appear at each end of the arrow to indicate that the connector has been selected.
2. Select **Relationship**.
3. Choose the new relationship you want to establish.
4. Save your changes.

## Deleting relationships in a job

At times, you may want to delete relationships between steps in a job. If you delete a step from a job, any relationships that were established between the deleted step and other steps in the job are automatically deleted.

If you simply want to modify a relationship, you do not need to delete it. See the topic “Viewing general relationship properties and modifying relationships” on page 125 for more information.

To delete a relationship between two steps:

1. In the Job Editor, right-click on the connector icon. Two small black boxes appear at each end of the arrow to indicate that the connector has been selected.
2. Select **Delete**. The relationship is deleted.
3. Save your changes.

---

## Saving jobs

It is recommended that you save changes to your jobs regularly. Changes that you make while working with jobs are not committed to the system until you save the job. In the Job Editor, a job containing unsaved changes is preceded by an asterisk in its name.

If you attempt to close a job before you have saved your changes, the system prompts you to save changes. The dialog box that opens depends on how you chose to close the job.

**Closing a job in the Job Editor.** If you attempt to close a job with unsaved changes in the Job Editor, a Save Resource dialog box opens.

You have the following options in this dialog box:

1. To save your changes and close the job, click **Yes**.
2. To disregard your changes and close the job, click **No**.
3. To exit the dialog box and return to the job, click **Cancel**. The Job Editor reopens.

**Exiting the system.** If you attempt to exit the system and you have unsaved jobs in the Job Editor, a Save Resources dialog box opens. This dialog box provides a list of the jobs with unsaved changes in them.

To save the changes:

1. Select the check boxes next to the jobs with changes you want to save. To expedite the process, click **Select All** to save changes in all of the jobs in the list, or click **Deselect All** to discard changes in all of the jobs in the list.
2. Click **OK**.

To exit the system without saving changes to any of the jobs in the list, click **Cancel**.

It is important to note that if any changes are made to the files in a job—for example, an IBM SPSS Modeler stream (*.str*) or a scenario file (*.scn*), any job that contains the file is affected. When changes are made to the file, a new version of the file is saved to the repository. However, the job that contains the file is not automatically updated with the modified file. To incorporate the file updates into the affected job:

1. Reopen the job. When the job is reopened, an asterisk appears with the job name in the job canvas, indicating that the job contains unsaved changes.
2. Resave the job.

---

## Job step results

Most types of IBM SPSS Collaboration and Deployment Services Deployment Manager job steps generate results in the form of output files. The files can be saved in the IBM SPSS Collaboration and Deployment Services Repository or in the file system running the application that generated them (for example, a system running IBM SPSS Statistics server). Certain types of job steps can generate multiple output files. Most types of job steps allow you to select output formats, such as HTML, text, or PDF.

Job step output settings are displayed on the Results tab of the Job Editor for IBM SPSS Statistics, SAS, IBM SPSS Modeler, and reporting job steps and on the Output Files tab for general job steps. The Results tab and available options are different for different job step types.

Regardless of the job step type, the following output file settings can be modified:

- Output file location
- Output file permissions
- Output file metadata

## Output file location

Click the ellipsis button in a selected Location cell to define the location for the results.

**Save to repository.** Saves output files to the repository in the specified folder. Click **Browse** to select the folder in which to save the output. If the files already exist, new versions are saved.

**Save to the application server file system.** Saves output files to the application server running the IBM SPSS Collaboration and Deployment Services Repository. In the Folder field, type the name of the application server folder in which to save the file.

**Discard.** Job output is discarded.

## Output file permissions

When a job is run, any output generated by the job is owned by the user that ran the job. Output file permissions for additional users or groups can be set in the Output Permissions dialog box. This dialog can be used to:

- Add a new principal. See the topic “Adding new users or groups” on page 28 for more information.
- Add a principal using a variable. See the topic “Add principal by variable” on page 128 for more information.
- Delete an existing principal. See the topic “Deleting a user or group from the permissions list” on page 29 for more information.
- Modify permissions for an existing principal.

Accessing the Output Permissions dialog box varies by the job step type.

- **IBM SPSS Statistics, SAS, and IBM SPSS Modeler job step output files.** In the Results tab, click the ellipsis in the Permissions column next to the file entry.
- **Reporting job step output files.** Click **Browse** next to the Permissions field.
- **General job step output files.** On the Output Files tab, click the ellipsis in the Permissions column next to the file entry.

To modify permissions for an existing user or group:

1. Select the principal from the table.
2. In the Permissions column, click the drop-down arrow and select a new permissions level from the list.
3. Click **OK**.

## Add principal by variable

Job step output file permissions can be assigned dynamically by adding users and groups to the permissions list as iterative variable values. In such cases, the variables defined within the step are used to retrieve the recipient address value every time the step is executed.

See the topic “Type for reporting steps” on page 158 for more information.

To add a user or group to the permissions list, in the Output Permissions dialog box click **Add Variable**. The Add Principal By Variable dialog box opens.

1. Select the security provider for authenticating the principal defined by the iterative variable.
2. Select **User** or **Group** as the principal type.
3. Enter the variable in the Principal Variable field. Use field content assist if necessary. See the topic “Entry field content assist” on page 10 for more information.

## Output file metadata

IBM SPSS Collaboration and Deployment Services Deployment Manager provides the capability to specify metadata properties for IBM SPSS Statistics, IBM SPSS Modeler, SAS, and reporting job step output files similar to the way properties are specified for IBM SPSS Collaboration and Deployment Services Repository objects.

To specify properties for a job step output file:

1. Click the step in an open job.
2. Click the Results tab.
3. In the list of output files, click in the Properties column next to the file. The ellipsis button is displayed. Disregard this step for reporting job steps.
4. Click the ellipsis button. The Output Properties dialog box opens. For reporting steps, click **Browse** next to the **Metadata** field.
5. Provide the following information.

- **Description.** A user-defined description of the output file.
- **Keywords.** The metadata assigned to the output file for the purposes of content search.
- **Author.** A string identifying the author of the output file.

*Note:* The dialog box may also display custom properties if such properties have been defined for the server. See the topic “Creating custom properties” on page 39 for more information.

- **Expiration Date.** The date after which the file is no longer active. By default, no expiration date is set for the job output file. You can specify a time interval from the current date for file expiration (for example, one year) or enter a specific date.
  - **Version Labels.** A user-defined label for the file. By default, no label is assigned to the file. To specify a new label or select from a list of existing labels, click **Browse**. The Edit Version Labels dialog appears. See the topic “Editing the version label” on page 30 for more information. If your label permissions change after defining output labels, labels may appear in the list that can no longer be applied. Job execution will fail due to an inability to assign a label. The label that cannot be assigned must be removed from the output properties or your permissions to the label will need to be changed. See the topic “Label security” on page 34 for more information.
  - **Topics.** The topics assigned to the file. Click **Add** to open the Add Topic dialog box and select topics. To remove a topic, select the topic in the list and click **Remove**. See the topic “Working with topics” on page 41 for more information.
6. To save your changes, click **OK**.

*Note:* Fields marked with a light bulb icon support the insertion of variables to the property at run time. For example, you can append the date to an output file/folder name at run time. For burst reporting job step types, you can also insert burst variables. See “Entry field content assist” on page 10.

---

## Chapter 11. Running jobs

Running a job results in all of its job steps executing according to their relationships defined in the job. Job execution can be initiated on-demand in response to a user request or scheduled to occur in the future in response to a specified event.

The job status and job history are displayed in the history view for jobs run through schedules and jobs run immediately. See the topic “Job history view” on page 137 for more information. Existing time-based and message-based schedules are listed in the job schedule view. See the topic “Job schedule view” on page 136 for more information.

---

### On-demand job execution

At times, you may want to run a job immediately after you have created or modified it. This process, which is typically used to test a job, is launched manually.

To run a job immediately:

1. In the Content Explorer, select an existing job.
2. From the **Tools** menu, select **Run Job Now**. The Job Execution dialog box opens. Alternatively, you can select the **Run Job Now** button from the toolbar.
3. Click **OK**.

The execution uses the specified default values for any variables in the job. In addition, any notifications associated with the job and its steps will be processed and delivered. To specify alternative variable values or disable notifications for the execution, use **Run Job with Options**.

It is not necessary to open a job in order to run it. You only need to select the job. If you want to view the contents of the job, double-click the job to open the Job Editor. Job execution can also be initiated from the Job Editor and the Job History view. See the topic “Job history view” on page 137 for more information.

### Specifying on-demand execution options

The default variable values and notification treatment can be overridden by specifying execution options. For example, it may be desirable to turn off notifications when testing jobs.

To specify options when running a job on-demand:

1. In the Content Explorer, select an existing job. See the topic “Opening an existing job” on page 120 for more information.
2. From the **Tools** menu, select **Run Job with Options**. The Run Job with Options dialog box opens.
3. Modify variable values as needed by clicking the value to be changed and typing the new value.
4. Specify whether or not the job and job step notifications should be delivered as a result of the execution.
5. Click **OK**.

The job executes with the specified options.

---

## Scheduled job execution

A job often needs to execute multiple times in response to future events. For example, a report created by a job may be needed every Friday. To define the execution pattern, the job must be associated with a schedule specifying the execution settings. A schedule can be either time-based or message-based.

- A time-based schedule relies on the occurrence of a specified time or date to initiate job execution. For example, a time-based schedule may run a job at 5:00 PM every Thursday.
- A message-based schedule relies on the receipt of a JMS message to trigger job execution. In this case, the job executes every time a specified message domain receives a message.

Given that a scheduled job executes at some point in the future, the schedule must use credentials defined in the system. Jobs cannot be scheduled using single sign-on (SSO) credentials because a user might not be logged in when the job executes.

## Creating schedules

To create a new schedule for a job, right-click on a job in the Content Explorer and select:

**New Schedule > Time-Based**

or

**New Schedule > Message-Based**

The Job Schedule wizard opens. Schedules for jobs must specify the following information:

- Settings for the job being scheduled
- Settings for the schedule that depend on whether the schedule is time-based or message-based
- Values for job variables

### Job information

The Job Information page of the Schedule wizard identifies the version of the job associated with the schedule and the credentials used for job execution.

**Job.** The name and IBM SPSS Collaboration and Deployment Services Repository path for the job associated with the schedule.

**Credentials.** Credentials, which specify permission levels, determine the user under which the scheduled job runs. Specifically, the system verifies credentials when writing output and saving files to the IBM SPSS Collaboration and Deployment Services Repository. Click **Browse** to select a credential from those defined in the system. To create a new credential for the scheduled job, click **New**. The specification of a credential for a job schedule is required. In Active Directory or OpenLDAP-based single sign-on environment, Server Process Credential can be used instead of regular user credential. See the topic “Server Process Credential” on page 48 for more information.

*Note:* If the job step is run under Active Directory user credentials, the credentials definition must be associated with the corresponding Active Directory domain. See the topic “Credential destination” on page 47 for more information.

**Label.** The label for the version of the job being scheduled. Select the value from the list of existing labels. The user under which the job runs must have adequate authority to the selected label when the schedule triggers job execution. If the user authority does not permit referencing the job by the selected label when the schedule triggers, the system will be unable to find the job and the execution will fail. See the topic “Label security” on page 34 for more information.



## Time-based schedule settings

For time-based schedules, the Schedule Time and Recurrence page of the Schedule wizard defines the time of day and frequency at which the job runs, as well as the duration of the schedule.

**Start Time.** Select the time of day at which you want to run the job from the drop-down list. To schedule a job at a time that is not available from the drop-down list (for example, 5:45 AM), type the value in the **Start time** field.

**Recurrence Pattern.** The recurrence pattern defines how often the job runs. Select the pattern and frequency from the following options:

- **Once.** The job runs one time only. From the **Date** drop-down list, select the date on which you want the job to run.
- **Hourly.** The job runs at hourly intervals. Specify the schedule frequency using the **Recur every <value> hours** field. For example, a value of 2 results in the job running every two hours.
- **Daily.** The job runs regularly at daily intervals. Specify the schedule frequency using the **Recur every <value> days** field. For example, a value of 3 results in the job running on every third day.
- **Weekly.** The job runs at weekly intervals on the specified day. Specify the schedule frequency using the **Recur every <value> weeks** field. For example, a value of 4 results in the job running on every four weeks on the designated day.
- **Monthly.** The job runs at monthly intervals on the specified day of the month. Specify the schedule frequency using the **Recur every <value> months** field. For example, a value of 4 results in the job running every fourth month on the designated day.

**Range of recurrence.** The range of recurrence is the duration for which the job runs and consists of two parts:

- **Start date.** The date for the initial execution of the job.

**Note:** The earliest that a daily schedule can begin is the day after it is created.

- **End date.** The date for the final execution of the job. For daily schedules, specify a date one day after the last execution day to allow the schedule to execute at the required time on the final day. Select **No end date** to have the job run indefinitely according to the specified schedule settings.

Click **Next** to specify values for variables used in the scheduled job.

## Message-based schedule settings

A message-based schedule is triggered by an outside event signaled by a JMS (Java Messaging Service) message. For example, when IBM SPSS Collaboration and Deployment Services job relies on the input from a third-party application, the application must send IBM SPSS Collaboration and Deployment Services a JMS message when the input file is ready for processing. For message-based schedules, the Message-Based page of the Schedule wizard defines the message domain and filter for the schedule.

**Message Domain.** The message domain identifies the JMS topic to which to subscribe. Select the domain from the list or click **New** to create a new message domain. See the topic “Message domains” on page 54 for more information.

**Message Filter.** Optional values that must be satisfied by a message for schedule activation to occur. Filters can be based on the text and the header of the message.

- **Message Text.** For JMS text messages, text that must be included in the message to activate the schedule.
- **Message Selector.** Optional selector text of contents of the message header. For example, you can indicate that the message header must contain the ID of a specific repository resource (ResourceID=<resource ID>) or specific custom properties (NewsType='Sports' or NewsType='Business').

**Use Durable Subscriptions.** The option to save the messages specified for the schedule while IBM SPSS Collaboration and Deployment Services is not actively listening on the messaging service, so that the saved message can be retrieved later when the system starts listening.

*Note:* Job processing may depend on several different external events. The initial event that triggers the start of the execution is specified in the message-based schedule. If processing requires subsequent events to take place, they must be specified by message-based job steps. See the topic Chapter 18, “Message-based job steps,” on page 183 for more information.

Click **Next** to specify values for variables used in the scheduled job.

## Job variables for schedules

The Job Variables page for the Schedule wizard specifies values for the variables defined for the job associated with the schedule.

The variables table contains the following columns:

- **Name.** The names of the existing job variables.
- **Value.** The value currently assigned to each variable. These are typically the default values defined for the variables.
- **Description.** Informative text about each variable.

When scheduling a job, you may want to use alternative values for one or more variables. To change a variable value:

1. Click the value to be changed.
2. Type the new value for the variable.
3. Press the Enter key.

The existing value updates to the new value. To return the variables to their default values, click **Restore Defaults**.

Click **Finish** to create the schedule for the job.

## Mapping JMS header properties to job variables:

For message-based schedules, the JMS message may contain variables in the header that can be mapped to the values for job variables. This mapping may be manual or automatic, as defined by the Job Variables page of the Schedule wizard.

In manual mapping, the header properties from the message are assigned to specific job variables. Message header variables are referenced using the following syntax:

```
#{JMSHeader.propertyName}
```

The value of *propertyName* corresponds to the name of the property in the message header. For example, suppose the header for the message associated with the schedule includes the property *SalesRegion* and the job being scheduled includes the job variable *region*. The value of *SalesRegion* is assigned to *region* by specifying `#{JMSHeader.SalesRegion}` as the value for *region*.

In contrast, for automatic mapping, the values for header properties in the JMS message are automatically used for any job variables that have the same name as the header property. Values defined in the schedule for the matching job variables are replaced with values from the message. Job variables that do not match a header property use the values specified in the schedule. If any of those values reference a header property that is unavailable, the name of the value is used. To use an empty string instead of the value name, use **quiet reference notation** by inserting an exclamation point between the dollar sign and the property name:

```
${!JMSHeader.propertyName}
```

Header properties in the message that do not match a job variable are unused by the schedule. To enable automatic mapping, select the **Automatically map JMS header properties to job variables** option.

## Editing schedules

A list of schedules defined in the system can be accessed using the Job Schedule view. See the topic “Job schedule view” on page 136 for more information.

To edit an existing job schedule, right-click on the schedule in the Job Schedule view and select **Edit Schedule**. Alternatively, you can click the *Edit Schedule* icon.

The Schedule wizard opens, showing the settings for the selected schedule. Modify the values as needed and click **Finish** to save the updated schedule.

## Reactivating a dormant schedule

If a version label of a job is removed, or if that version of the job is deleted, the schedule associated with the labeled version becomes dormant. Thus, the dormant schedule cannot be used until it is reassociated with a valid labeled job version.

If a schedule becomes dormant, the following message appears:

The scheduled label, <label name>, for this job does not exist anymore. You must select a different label, or re-apply the originally scheduled label to the job.

<label name> represents the version label that was removed or the job version that was deleted.

The resolution for a dormant schedule depends upon how the schedule became dormant.

- **Deleted job version.** If the job version was deleted, a different job version must be selected for scheduling.
- **Deleted version label.** If the label for the job version was removed, the label needs to be reapplied to the job version.

## Deleting schedules

To delete an existing job schedule, right-click on the schedule in the Job Schedule view and select **Delete Schedule(s)**. Alternatively, you can click the *Delete Schedule(s)* icon.

The job associated with the schedule remains in the system. However, the job will no longer execute in accordance with the deleted schedule.

## 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 12. Monitoring status

In IBM SPSS Collaboration and Deployment Services Deployment Manager, the status of a job can be analyzed via a number of job summary views. Information is organized in a table and is designed to provide an at-a-glance summary for jobs in the repository.

The following views are available:

- Job Schedule
- Job History
- Model Management
- Predictors

Although individual tables contain specialized information, all job summary views are accessed in a similar way.

In addition, the status of servers and server clusters can also be monitored. See the topic “Server status view” on page 144 for more information.

---

### Accessing status views

In the following section, the <View Type> refers to the specialized view. Examples include:

- *Job Schedule*
- *Job History*
- *Model Management*
- *Server Status*

Job summary views can be accessed using any of the following methods:

**Toolbar.** This option launches a blank job summary view. To access the job schedule view from the toolbar:

1. From the **View** menu, select:  
**Show View > <View Type>**

**Content Explorer.** This option launches a populated job summary view. To access the job summary view from the Content Explorer:

2. Right-click on a job and select **Show <View Type>**.

---

### Selecting a server in status views

By default, jobs on the current server appear in the job list.

To view jobs for other servers:

1. From the **Server** drop-down list, select a server.

The job list is updated with jobs on the selected server.

*Note:* If you are not logged into the server that was selected, the *Login to IBM SPSS Collaboration and Deployment Services Repository* dialog box opens, prompting you to log into the server.

---

## Opening a job in the Job Editor

At times, when a job summary table is open (e.g., the job schedule or the job history), it may be useful to open the job in the Job Editor and view the job contents in conjunction with the job table.

To open a job in the Job Editor:

1. Select the job in the table.
2. Click the **Open job in editor** icon. The job opens in the canvas of the Job Editor.

---

## Refreshing status views

When a job is run, the corresponding job tables (e.g., job schedule and job history) are not refreshed automatically. You must refresh job tables manually.

To refresh the job table:

1. Click the Refresh icon. The job table displays the updated status.

---

## Reordering items in status views

By default, items in the job summary tables (e.g., job schedule and job history) are organized in chronological order. The order of items in the job summary tables can be reorganized by column.

To reorder items:

1. Click in the header row of the column that you want to reorder.
2. Click on the arrow in the row title. The table is reordered.

---

## Deleting a job from status views

To delete an item from a job summary table (e.g., a job schedule or a job history):

1. Select the item in the table. To select multiple items, hold down the **Ctrl** key and select additional rows.
2. Click the **Delete** icon. The Delete Confirmation dialog box opens.
3. Click **OK**. The item is removed from the job schedule table.

---

## Job schedule view

A scheduled job is run automatically at a designated date and time or when a JMS message is received from an outside application. Job schedules are described in the Job Schedule table, which contains the following summary information:

**Job Name.** The name of the job.

**Version Label.** The label that has been applied to the job. If no label has been applied to the job, this field is blank.

**Schedule Summary.** The frequency with which the job is scheduled to run. Valid values are *Once*, *Hourly*, *Daily*, *Weekly*, and *Monthly*.

**Next Start.** The next date and time at which the job is scheduled to run.

**Last Run.** The date and time when the job was last run.

**Last Run Status.** The status of the job when it was last run. Valid values include *SUCCESS* or *FAILED*. A canceled job is recorded as a failure.

**Credential.** The credential under which the scheduled job will run.

The job schedule type (message-based or time-based) is specified by the icon in the column before the job name.

The first column indicates any problems with a schedule, such as a reference to a message domain that has been deleted, by displaying a warning icon. Right-click a problematic schedule and select **Show Error Messages** to view any errors associated with the schedule.

---

## Job history view

Each time a job is run, the action is recorded in the job history, which provides status information about the job and its corresponding steps.

Specifically, the *Name* column describes the overall job. When the + is expanded, information about the individual job steps within a job appears.

Information in the job history table cannot be modified because this table reflects information obtained from other components in the system. Job information is provided on a per server basis. See the topic “Selecting a server in status views” on page 135 for more information.

In addition, filters can be applied to the job history to reduce the number of jobs that appear in the list. See the topic “Job History Filters” on page 143 for more information.

## Working with the job history table

The job history table contains the following information:

**Name.** The name of the job or step.

**Version.** The version label that has been applied to the job. If no version label was explicitly applied in the job schedule, then the *LATEST* label is used by default.

**Status.** The current status of the job or step. Valid values include *Success*, *Running* or *Failed*. A canceled job is recorded as a failure. The status of the individual job steps that comprise each job, as well as any corresponding logs, appear underneath each job. To expand the job history list for a specific job, click the + next to the job.

**Start date.** The date and time at which the job or step was started.

**Run time.** The time taken to run the job or step. It is important to note that a value in this field does not necessarily imply success. The *Status* column describes whether or not the job ran successfully.

**User.** The user who last scheduled the job.

### Blank cells in the job history view

If no label has been applied to a job, the corresponding **Label** field is blank. For the **Status**, **Start Time**, and **Run Time** fields, a blank cell indicates that the job has not been run yet.

### Canceling a job

You can cancel a job while it is running. It is important to note that canceling a job is not the same as deleting a job. Canceling only stops the job from running. Canceling does not remove a job from the content repository.

**Important:** You must be assigned the *Schedule* action to be able to cancel jobs.

If you cancel a job that contains relationships, the relationships will not be followed. For example, suppose your job contains step A and step B, which are joined by a sequential connector. If you cancel the job while step A is running, step B will not be executed. This process applies to all relationship connectors, including the fail connector.

To cancel a job:

1. In the Job History table, select the job that you want to cancel. The job status of the selected job should be *Running*.
2. Click the **Stop selected running job** icon.

## Viewing job results

If results are available, separate rows appear under the job steps in the table. Each row contains the path in which the results are stored. For certain applications, such as IBM SPSS Modeler, double-clicking on the results path launches the IBM SPSS Modeler client and displays the results.

The system can open any result file that is written back to the IBM SPSS Collaboration and Deployment Services Repository. If the file type has a corresponding editor in the IBM SPSS Collaboration and Deployment Services Deployment Manager (e.g., HTML, text, image files), the system opens the file within the IBM SPSS Collaboration and Deployment Services Deployment Manager. If a file path is preceded by the machine name, the IBM SPSS Collaboration and Deployment Services Deployment Manager does not have an editor to display the file and will call upon the operating system to open the file.

## Viewing job logs

The job history table includes logs for the overall job as well as for steps within the job. These logs are system-generated logs and cannot be modified. Note that some types of steps do not produce logs.

In the table, the overall log appears under a subheading corresponding to the name of the job. This log contains information relevant to the entire job, such as whether or not residual artifacts have been deleted. The values used for any job variables also appear in this log.

If a log is available for a job step, it too can be accessed from the job history table. Job step logs appear under a subheading corresponding to the job step name and pertain to the corresponding step only. If the job contains multiple steps that produce logs, the logs appear under each step in the table. These are not logs for the entire job, and they are not logs for the system.

To view a log, double-click **Log** in the Name column for the cell that corresponds to the log that you want to view. The log appears in a separate editor.

---

## Model management views

A summary of model results is provided in the model management views. Similar to the job history and job schedule views, the model management tables provide information about model analyses.

The following views are available:

- Model evaluation
- Champion challenger

The information provided in each view, which may vary by model type, is not modifiable. In addition, filters can be used to narrow the results of a model management view even further. See the topic “Model management filters” on page 144 for more information.

To select a specific model management view:

1. From the **View** menu select:  
**Show View > Model Management**



The Model Management tab appears.

2. From the **Server** drop-down list, select the server name.
3. From the **Type** drop-down list, select the view type. Options include model evaluation or champion challenger.

**Note:** The model management views present information about IBM SPSS Modeler files stored in the repository. To store these files in the repository, the IBM SPSS Collaboration and Deployment Services environment must include IBM SPSS Modeler adapters. For information on installing the adapters, see the IBM SPSS Modeler documentation.

## Model evaluation view

The model evaluation view describes results for model evaluation scoring branches. This view is designed to provide an at-a-glance summary of the overall performance of the branch. For example, this view describes whether the branch is trending upward, downward, or has remained the same. Jobs may appear multiple times in the model evaluation view. One entry appears for each labeled version of a job that contains a model evaluation scoring branch.

The Model Evaluation view contains the following information.

- **Scoring Branch.** The branch containing the scoring node.
- **File.** The file containing the scoring branch.
- **Version.** The version of the file used to generate the results.
- **Index.** A percentage value. The colored circle next to the index value corresponds to the thresholds for performance (e.g., good, better, bad) that were specified in the General tab for the job step. For example, a red circle indicates that the index value falls into the lowest-performing range.
- **Trend.** The percentage change in the model. An arrow indicates whether the trend is up or down. If the trend is 0.00, a horizontal bar appears, which represents no change. The first time the file is run, the trend field is blank.
- **Author.** The author of the file.
- **Type.** The type of analysis. Examples include *gains* and *accuracy*.
- **Data.** The actual data used for the source node. The value depends on the source node type. For example, for an Enterprise View node, the value corresponds to a Data Provider Definition. For an ODBC node, the value is a DSN name. For a Variable File node, the value is a file name.
- **Job.** The job that references the file.
- **Job Version.** The version of the job used to analyze the model.
- **Last successful run.** The most recent date and time at which the job was run successfully. All of the model evaluation views only display information for the most recent job that was run successfully. A full job history, including any previously run jobs or failed jobs is available in the job history view. See the topic “Job history view” on page 137 for more information.

## Champion challenger view

The Champion Challenger view describes files that were compared to one another within a Champion Challenger job step. The file containing the scoring branch that was deemed to be the most effective is designated as the champion. The champion file appears as the uppermost row in the Champion Challenger view. To see all of the branches analyzed in the Champion Challenger job step, simply expand the tree.

It is important to note that one of the scoring branches in the list was designated the champion. The champion file is then renamed in accordance with the parameters specified on the Champion tab of the champion-challenger job step. Thus, the renamed file appears in the uppermost row.

The Champion Challenger table contains the following information.

- **Champion.** Champion-challenger files are organized around the champion. Thus, the first line introducing a set of champion-challenger files contains the name and path information of the champion file. The files underneath the champions are challenger files.
- **Is Copy.** This column indicates whether or not the champion file listed in the first row of the current set of champion-challenger files is a copy of the winning file. (The winning file is designated by a star in the Index column.) If *Yes* appears in this field, then a copy of the champion file is displayed in the first row. By default, IBM SPSS Collaboration and Deployment Services Deployment Manager makes a copy of the winning champion file. This setting can be disabled when creating or editing a champion-challenger job step. If the option to copy the champion file was cleared, then the winning champion file appears in the first row under the scenario name.
- **File.** The file containing the scoring branch.
- **Version.** The version of the file used to generate the results.
- **Index.** A percentage value. If a star appears next in the Index column, this file was the winning file in the champion-challenger analysis.
- **Trend.** The percentage change in the model. An arrow indicates whether the trend is up or down. If the trend is 0.00, a horizontal bar appears, which represents no change. The first time the file is run, the trend field is blank.
- **Author.** The author of the file.
- **Type.** The type of analysis. Examples include *gains* and *accuracy*.
- **Source.** The name of the source node that supplies the data for the scoring branch. This may not be the original source node defined in the file. For Champion Challenger steps, a source node from any of the challengers may be used. The column values include the name of the file containing the source node as a prefix before the node name to identify the location of the source node used.
- **Data.** The actual data used for the source node. The value depends on the source node type. For example, for an Enterprise View node, the value corresponds to a Data Provider Definition. For an ODBC node, the value is a DSN name. For a Variable File node, the value is a file name.
- **Job.** The job that references the file.
- **Job Version.** The version of the job used to analyze the model.
- **Job step.** The name of the job step that references the file.
- **Last successful run.** The most recent date and time at which the job was run successfully. All of the model evaluation views only display information for the most recent job that was run successfully. A full job history, including any previously run jobs or failed jobs is available in the job history view. See the topic “Job history view” on page 137 for more information.

## Predictors view

Both champion challenger and model evaluation files use predictors to generate results. A predictor is a variable that is used as an input for a model. Typically, a model contains multiple predictors. Predictors are evaluated and then ranked based on their relevance to the results.

To view the Predictors table:

1. In the Content Explorer, select an object.
2. From the **View** menu select:

### Show View > Predictors

The Predictors table appears.

The predictor effectiveness table contains the following information. Some parameters apply at the predictor level. Other columns in the table apply to the entire file.

## Predictor-level information

**Name.** The predictor name. Predictors are organized by file. Thus, the first line introducing a set of predictors contains the file name and path information.

**Source.** The data source used to obtain predictors—for example, *Enterprise View*.

**Value.** The value of the predictor.

**Rank.** The rank of the predictor. Predictors are listed in order of descending rank. The rank of a predictor corresponds to its importance within the model. For example, if the a predictor for household income is listed first on the list, then for this model, the level of household income was most highly correlated with a positive response.

## File-level information

**Version.** The version of the file used to generate the results.

**Author.** The author of the file.

**Data provider definition.** The data provider definition (DPD) associated with the file.

**Application view.** The application view associated with the file.

**Job.** The job that references the file.

**Application view label.** The label associated with the application view used by the file.

**Last successful run.** The most recent date and time at which the job was run successfully. All of the model evaluation views only display information for the most recent job that was run successfully. A full job history, including any previously run jobs or failed jobs is available in the job history view. See the topic “Job history view” on page 137 for more information.

**Note:** The Predictors view presents information about IBM SPSS Modeler files stored in the repository. To store these files in the repository, the IBM SPSS Collaboration and Deployment Services environment must include IBM SPSS Modeler adapters. For information on installing the adapters, see the IBM SPSS Modeler documentation.

---

## Filters

Any of the available views can be filtered to reduce the number of results that appear in the view table. Although some filtering options are common to all views, specific filtering parameters may vary by the view type. Multiple filters can be used simultaneously. For example, the job filter is often used in conjunction with other filters.

By default, filtering is disabled. Once filtering is enabled, the ability to filter the job history persists across IBM SPSS Collaboration and Deployment Services Deployment Manager sessions. In addition, filter settings persist across server connections. For example, if filtering is enabled and the server selection is changed from server A to server B, the filter settings established for server A persist for server B.

## Filters common to all status views

The following filters are common to all views:

**Job.** The jobs that appear in the view table can be limited to the following:

- **Selected job in job editor.** Filters the table, so that only information for the currently selected job appears.
- **User-defined job.** Used to search for a job schedule by name. Selecting **Browse** allows a search of the entire content repository.

**Version Label.** Limits the list to objects containing the specified version label. Typically, this option is used in conjunction with the another filter. For example, if the version label *LATEST* were chosen without a corresponding job filter, every job in repository would appear in the job history because the current version of any job is implicitly labeled *LATEST*. For the model management filters, the results can be limited by job version label or file version label.

## Job schedule filters

To access job schedule filters, click the **Filters** button in the job schedule view.

**Enable Filtering.** If filtering has not been previously enabled, select the **Enable Filtering** check box.

The job schedule can be filtered based on any of the following criteria:

**Job.** The jobs that appear in the job schedule can be limited to the following:

- **Selected job in job editor.** Filters the job schedule table, so that only the history for the currently selected job appears.
- **User-defined job.** Used to search for a job schedule by name. Selecting **Browse** allows a search of the entire content repository.

**Version Label.** Narrows the list to jobs containing the specified version label. Typically, this option is used in conjunction with the job filter. For example, if the version label *LATEST* were chosen without a corresponding job filter, every job in the repository would appear in the job schedule because the current version of any job is implicitly labeled *LATEST*.

**Last Run Status.** Narrows the list to jobs containing the selected status. The following status options can be specified:

- Success
- Failure
- Canceled

**Schedule type** Enables filtering by schedule type (message-based or time-based).

- **Time-Based Fields.** Selecting the option allows you to specify the criteria for filtering time-based schedules. This limits the view to schedules having a next start time within the specified time interval.

Table 12. Time-based filtering options

Field	Description
Relative Time Period	A finite time period relative to the current date. The following options are available: <ul style="list-style-type: none"> <li>• Current day</li> <li>• Current week</li> <li>• Current month</li> <li>• Current quarter</li> <li>• Current year</li> <li>• Most recent 30 days</li> <li>• Most recent 3 months</li> <li>• Most recent 6 months</li> <li>• Most recent 9 months</li> </ul>
Absolute Time Period	A specific date-based range of time. A valid range must be provided--for example, the end date must be equal to or greater than the start date.

- **Message-Based Fields.** Selecting the option allows you to specify the following criteria for filtering message-based schedules:

Table 13. Message-based filtering options

Field	Description
Message Domain	Filters schedules based on the JMS topic of the corresponding messages.
Message Text	Filters schedules based on the text of the corresponding JMS text messages.
Message Selector	Filters schedules based on the contents of the message headers, for example <code>NewsType='Sports'</code> or <code>NewsType='Business'</code>

## Job History Filters

To access job history filters, click the **Filters** button in the job history view.

**Enable filtering.** If filtering has not been previously enabled, select the **Enable Filtering** check box.

The job history can be filtered based on any of the following criteria:

**Job.** The jobs that appear in the job history can be limited to the following:

- **The selected job in job editor.** Filters the job history table, so that only the history for the currently selected job appears.
- **A user-defined job.** Used to search for a job history by name. Selecting **Browse** allows a search of the entire content repository.

**Version Label.** Narrows the list to jobs containing the specified version label. Typically, this option is used in conjunction with the job filter. For example, if the version label *LATEST* were chosen without a corresponding job filter, every job in repository would appear in the job history because the current version of any job is implicitly labeled *LATEST*.

**Last Run Status.** Narrows the list to jobs containing the selected status. The following status options can be specified:

- Success
- Failure
- Canceled

**Mode of execution.** Narrows the list by mode of execution. Valid values include:

- Scheduled—Jobs that were previously scheduled and ran automatically at the specified date and time.
- Manual—Jobs that were executed manually using **Run Job Now**.
- Submitted—Jobs that were submitted for execution by an external source.

**Start Date.** Narrows the job history list based on the date the job was started. A relative or absolute time period can be specified.

- **Relative time period.** A finite time period relative to the current date. The following options are available:
  - Current day
  - Current week
  - Current month
  - Current quarter
  - Current year
  - Most recent 30 days
  - Most recent 3 months
  - Most recent 6 months
  - Most recent 9 months
- **Absolute time period.** A specific date-based range of time. A valid range must be provided—for example, the end date must be equal or greater than the start date.

## Model management filters

The results that appear in the model management view can be filtered based upon the following criteria.

**Trend.** The trends that appear in model management views can be limited to the following:

- **Direction.** Shows results having the same trend direction. For example, you can limit the view to all results trending up.
- **Values.** The range of the percentage change. A minimum and maximum percentage must be specified.

**File.** Limits the results to scoring branches contained in a specific file.

**Index.** A percentage value. A minimum and maximum percentage must be specified.

**Evaluation type.** Valid values include *Gains*, *Accuracy*, and *Accreditation*.

**Data.** The data used for the source node

**Model Execution Date.** A range of execution dates for scoring branches. Specify the start and end dates for the range.

---

## Server status view

The status of servers and server clusters is available in the Server Status view. This view provides metadata about servers and server clusters for which you have permissions.

Information cannot be modified in the server status table. Any modifications to server or server cluster configurations must be done in the *Resource Definitions* folder. See the topic Chapter 6, “Resource definitions,” on page 47 for more information.

To view the status for a server, from the **Server** drop-down list, select a server. A list of servers appears.

The server status table contains the following information:

**Host.** The host name for the server.

**Port.** The port number for the server.

**Type.** The type of server; examples include *IBM SPSS Statistics server*, *IBM SPSS Modeler server*, *Remote process server*, and *Cluster*.

**Up Since.** The last date and time at which the server was last started.

**State.** The status of the server—for example, *running*.

**Connections.** The number of connections into the server.

**Weight.** The weight of the server within a server cluster.

**Viewing Servers by Cluster.** By default, the server status table contains a list of servers. To view servers organized by cluster, click the show clusters icon. For server clusters, the host cell is blank. To view the servers within a cluster, click the + sign to expand the list. The host name for each server within the cluster appears in the table.

## Setting the refresh rate

By default, the server status view is not automatically refreshed. To modify the refresh rate:

1. From the **Refresh Rate** drop-down list, select the refresh rate increment—for example, *15 minutes*.





---

## Chapter 13. Notifications and subscriptions

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.

Notifications and subscriptions generate email messages, RSS feeds, or JMS (Java Messaging Service) messages when corresponding events occur. Notifications can be defined for multiple recipients while subscriptions to IBM SPSS Collaboration and Deployment Services Repository objects are defined by the system users only for themselves.

Notifications and subscriptions are retained when IBM SPSS Collaboration and Deployment Services Repository content is exported and imported.

---

### Notifications

Notifications can be defined for changes to the IBM SPSS Collaboration and Deployment Services Repository content and processing events.

File notifications are triggered by the creation of new file and job versions. Folder notifications are triggered by folder content changes, such as addition of files, creation of new file and job versions, and creation of subfolders. Already existing notifications can be modified or removed for a single object or for multiple repository objects. See the topic “Content notifications” on page 148 for more information.

A special kind of notifications are the notifications for content label events. These notifications are triggered when certain labels are applied to IBM SPSS Collaboration and Deployment Services Repository objects. See the topic “Label event notifications” on page 149 for more information.

Notifications can also be defined for job processing events such as job and job step success and failure. See “Job success and failure notifications,” “Job step notifications” on page 148 and “Notifications based on model evaluation return codes” on page 149 for more information.

A user must be assigned the corresponding action through roles to define notifications.

### Job success and failure notifications

IBM SPSS Collaboration and Deployment Services Repository provides the capability to send email notification messages for a job success or failure.

To display job notifications:

1. Open the job.
2. Click the **Notifications** tab. The Notifications for Success and Notifications for Failure fields display the number of recipients defined for each type of notification.

To display the list of notification recipients, place the mouse pointer over the Update button.

To specify or update job notifications:

1. On the **Notifications** tab, click **Update** next to the notification type. The Notifications dialog box appears.
2. Specify the notification sender and recipients, and customize the message as necessary. See the topic “Notification settings” on page 150 for more information.
3. Save the job.

**Note:** Notifications can also be defined for jobs as content repository objects. Such notifications will be triggered by the creation of a new version of a job. See the topic “Content notifications” for more information. For example, defining job success or failure notifications for a job or a job step creates a new version of the job and subsequently triggers notification for the change.

## Job step notifications

IBM SPSS Collaboration and Deployment Services Repository provides the capability to set up notifications for the success and failure of individual job steps, as well as jobs.

A special case of job step notifications are notifications of success and failure of individual iterations; such notifications can be defined only for iterative job steps.

To display job step notifications:

1. Open the job.
2. Select the job step.
3. Click the **Notifications** tab. The Notifications for Success and Notifications for Failure fields display the number of recipients defined for each type of notification.

To display the list of notification recipients, place the mouse pointer over the Update button.

To specify or update job step notifications:

1. On the **Notifications** tab, click **Update** next to the notification type. The Notifications dialog box appears.
2. Specify the notification sender and recipients, and customize the message as necessary. See the topic “Notification settings” on page 150 for more information.
3. Save the job.

## Content notifications

Content notifications keep users informed about changes to IBM SPSS Collaboration and Deployment Services Repository objects, such as the creation of new versions of files and jobs, and folder structures, such as the creation of new subfolders. Content notifications can be defined for a single object or multiple objects selected simultaneously.

To define or modify file or job notifications:

1. Right-click the file or job name and choose **Notifications**. To select multiple objects, hold down **Shift** or **Ctrl** key. The Notifications dialog box appears.
2. In the Notifications dialog box, specify the sender and recipients, customize the message as necessary, and, for folder notifications, indicate folder options. See the topic “Notification settings” on page 150 for more information.

To define or modify folder notifications:

1. Right-click the folder name and choose **Notifications**, then click **Folder Events** to define notifications for folder structure events or **Folder Content Events** to define notifications for changes to the objects in the folder.
2. In the Notifications dialog box, specify the sender, recipients, customize the message as necessary and, for folder notifications, indicate folder options. See the topic “Notification settings” on page 150 for more information.

To remove content notifications:

1. Right-click the file, job, or folder name and choose **Notifications**. The Notifications dialog box appears. Content notifications can also be removed for multiple objects selected simultaneously. To select multiple object, hold down **Shift** or **Ctrl** key.

2. In the Notifications dialog box, click **Remove Notification**. See the topic “Notification settings” on page 150 for more information.

Notes:

- The **Notifications** option is not available if objects of different kinds are selected simultaneously (for example, files and folders, folders and jobs).
- IBM SPSS Collaboration and Deployment Services jobs cannot be opened outside IBM SPSS Collaboration and Deployment Services Deployment Manager. If content notifications are defined for jobs, it is strongly recommended to modify the notification settings to remove the hyperlink to the job because it cannot be opened in a Web browser. See the topic “Customizing the notification message” on page 153 for more information.
- If you import files and folders that were previously exported, no notifications are sent for any import events.

## Notifications based on model evaluation return codes

Notifications can be set up to monitor specific model evaluation outcomes. The following are examples of when notifications might be based on model evaluation return codes:

- Notifications can be set up based on a specific outcome of the model evaluation. For example, if the result returned by a model evaluation analysis is less than 0.85, notifications could be sent to a predefined list of recipients.
- Notifications can be set up based on the traffic light level of a model evaluation job. For example, if the result is red, a notification could be sent. In addition, multiple notification rules could be set up with different email recipients for each traffic light level.

## Label event notifications

IBM SPSS Collaboration and Deployment Services Repository provides the capability to set up notifications for label events.

For example, an analyst can set up the *Production* label to notify the repository administrator when the label is applied to a job so that the job can subsequently be promoted to production status. Label event notifications can be limited to certain content types, for example, repository jobs or SPSS Statistics files.

To access label event notifications, in IBM SPSS Collaboration and Deployment Services Deployment Manager, right-click the Content Repository (root) folder and choose

**Notifications > Label Events...**

The Label Events dialog box appears. The dialog allows you to add, change, and delete the notifications for label events.

To add a label event notification, click the **Add a label event** button. Specify the notification sender and recipients, the label and content types, and customize the subject and body of the message as necessary. See the topic “Notification settings” on page 150 for more information.

To change a label event notification, select a label event notifications entry and click the **Change selected label event** button. Specify the notification sender and recipients, the label and content types, and customize the subject and body of the message as necessary. See the topic “Notification settings” on page 150 for more information.

To delete a label event notifications, select a label event notifications entry and click the **Delete selected label event** button.

## Notification settings

Notification settings, including the notification sender, recipients, message customizations, and folder options, are specified in the Notifications dialog box.

**From.** The address of the sender notification message. The field is prepopulated with the default email address in repository configuration options.

**To.** The list of notification recipients. The addresses can be entered, specified as an iterative variable for iterative job step notifications, specified as a dynamic list based on a query, selected from a directory listing of a supported email application such as Microsoft Outlook or Lotus Notes, or selected from available IBM SPSS Collaboration and Deployment Services users. To edit an address, click the ellipsis button next to it in the list. The list also specifies whether the modified object is to be attached to the message for each individual recipient. See the topic “Notification message attachments” on page 153 for more information. Recipients can also be specified as variables. To remove a recipient, click **Delete**.

**Subject.** The subject of the notification message. By default, the field is populated by the default template associated with the notification event. Modify the message subject, if necessary. See the topic “Customizing the notification message” on page 153 for more information.

**Message.** The body text of the notification message. The field is populated by the default template associated with the notification event. Modify the message body text, if necessary. See the topic “Customizing the notification message” on page 153 for more information.

**Remove Notification.** Click to remove notification. When multiple objects are selected, clicking this button will remove all notifications for the objects.

**Apply to subfolders.** For folder notifications, applies the subscription to subfolders.

**Label.** For label event notifications, the label triggering the notification.

**Notify when changes are made to these file types only.** For folder and label event notifications, limits notification to selected file types. For example, notifications can be set up only for repository jobs or IBM SPSS Statistics files. Click the ellipsis button to open the File Types dialog box. The File Types dialog opens. Click the file type entries to select or clear. Use the Ctrl or Shift keys to select multiple entries. After all file types have been selected, click **OK**.

**Preview.** Click to preview the notification message. See the topic “Notification message preview” on page 154 for more information.

### Note:

- When notifications are modified for multiple objects, the list of recipients also displays a number next to the email address indicating how many of the selected objects have the recipient in their notification settings. For multiple objects with different notification settings, the fields with varying values are blank and marked with alert signs. When notifications are modified for multiple folders with varying folder option settings, the options are marked with green squares.
- If you use the Security Subscriber option to select notification recipients, email notifications are not sent to members of IBM SPSS Collaboration and Deployment Services groups. You must select individual users instead of groups, or define a distribution list associated with the IBM SPSS Collaboration and Deployment Services group (and IBM SPSS Collaboration and Deployment Services security must be capable of retrieving the email address defined for the group).

## Entering notification recipient email addresses

To enter a recipient email address, on the Notification screen, click the ellipsis button next to the To field or click the down arrow next and then choose **Email Subscriber**. The Set email dialog box appears.

**Email Address.** The email address of the notification recipient. The address string must contain a @ character followed by a period.

1. Enter the email address. Only one address can be entered at a time.
2. Click **OK**. The address will appear in the To list of the Notifications dialog box.

### Selecting notification recipients from Microsoft Outlook

To select a notification address from Microsoft Outlook directory, in the Notifications dialog box, click the down arrow next to the To field and then choose **Outlook Subscriber**. The Select Email Addresses dialog box appears.

**Type Name or Select from List.** Filters the list of available users according to the string entered. Use this field to refine the list of available users by user name.

**Show Names from the.** The list of organizational units defined in the directory. Use this field to refine the list of available users by organizational unit.

**To-->.** Displays selected users.

**Advanced.** Advanced options include Find and Properties. Click **Find** for an Outlook advanced directory search. Click **Properties** to display detailed information for the selected Outlook entry.

1. To select a recipient, double-click a directory entry in the list. The address appears in the To--> field. To remove a selected recipient, select the entry in the To--> field and press Delete key.
2. After all recipients have been selected, click **OK**. The email addresses of selected recipients appear in the To list of the Notification for Job Success dialog box.

*Note:* If the **Outlook Subscriber** option is not available in the drop-down, verify that *CDO.DLL* is installed on your system (the default directory is *C:\Program Files\Common Files\System\MSMAPI\1033*). If the file is not present, the Collaboration Data Objects component of Microsoft Outlook must be installed from the Microsoft Office installation CD. The component can also be obtained from the Microsoft download site.

### Selecting notification recipients from Lotus Notes

To add a notification email address from Lotus, in the Notifications dialog box, click the down arrow next to the To field and then choose **Lotus Subscriber**. You will be prompted for the Lotus Domino Server credentials, and then the Lotus Subscribers dialog box appears.

1. To select a recipient, highlight a directory entry in the list and click **To ->**. The address appears in the To-> field. To remove a recipient, select the entry in the To-> field and click Remove. Click Remove All to clear all recipients.
2. After all recipients have been selected, click **OK**. The email addresses of the selected recipients appear in the To list of the Notification for Job Success dialog box.

### Selecting notification recipients from security subscribers

To send notifications to IBM SPSS Collaboration and Deployment Services users or groups, in the Notifications dialog box, click the down arrow next to the To field and then choose **Security Subscriber**. The Select User or Group dialog box appears.

1. From the **Select provider** drop-down list, select the entity that contains your user and group information.
2. In the Search field, type the first few letters of the user ID or group you want to add. To search for all available user IDs and groups, leave this field blank.
3. Click **Search**. The users and groups that correspond to your search appear in the dialog box.
4. Select one or more users or groups from the list.
5. Click **OK**. The users or groups appear in the To field of the Notifications dialog box.

To receive email notifications, users must specify a valid email address in their user preferences or on the supported system (LDAP, for example).

**Note:** If you use the Security Subscriber option to select notification recipients, email notifications are not sent to members of IBM SPSS Collaboration and Deployment Services groups. You must select individual users instead of groups, or define a distribution list associated with the IBM SPSS Collaboration and Deployment Services group (and IBM SPSS Collaboration and Deployment Services security must be capable of retrieving the email address defined for the group).

Users can also receive syndicated feeds (RSS/Atom feeds) if they have been granted the *Access Syndicated Feeds* action by an administrator. The feeds will be aggregated individually for the given security principal, based on individual subscriptions (such as file subscriptions) and notifications (such as job completion events and folder content events) created by administrators for this principal or any security group the principal belongs to. The user must be authenticated to access the feed. If authentication is successful and the user has been granted the *Access Syndicated Feeds* action, IBM SPSS Collaboration and Deployment Services will aggregate the feed based on the security principal identifier and security groups.

For more information about RSS feeds, see the IBM SPSS Collaboration and Deployment Services Deployment Portal help.

### Dynamic notification lists

A list of notification recipients can also be generated dynamically using a database query. The input for the notification list can be the results of a BIRT Report Designer for IBM SPSS report (\*.rptdesign).

To generate a list of notification email addresses dynamically, in the Notification for Job Success dialog box, click the down arrow next to the To field and then choose **Dynamic List**. The Dynamic List dialog box appears.

1. To select the file that will be used to dynamically retrieve the list of recipients, click **Browse** next to the Source query field and select the file from the Content Repository. Then select the labeled version to use.
2. Select the data set, data source, credentials, and the column containing the list of email addresses to retrieve from the source query:
  - **Data set** is the name of the database.
  - **Data source** is the URL of the server.
  - **Credentials** come from IBM SPSS Collaboration and Deployment Services resource definitions and must provide access to the specified data set and data source.
  - **Column with email** is the database column containing the address value retrieved by the query.
3. After all recipients have been selected, click **OK**. The source file name, version, data set, credentials and column appear in the To list of the Notification for Job Success dialog box.

### JMS subscriber

Notification events can generate JMS (Java Messaging Service) messages that trigger internal IBM SPSS Collaboration and Deployment Services processing as well as external applications.

For example, a IBM SPSS Collaboration and Deployment Services user can set up JMS notification for a IBM SPSS Modeler stream to rerun a job which includes the stream when a new version of the stream is created. For that, a JMS message domain must first be created. Then a message-based schedule must be set up for the job specifying the resource ID of the stream in the message selector. See the topic "Message-based schedule settings" on page 131 for more information. And finally, a notification must be set up for the IBM SPSS Modeler stream based on the JMS subscriber to be triggered when a new version of the stream is created.

To specify a JMS subscriber, in the Notifications dialog box click the down arrow next to the To field and then choose **JMS Subscriber**. The JMS Subscriber dialog box appears.

**Message Domain.** The JMS message domain for the subscriber. See the topic “Message domains” on page 54 for more information.

1. Select the message domain. A message domain can only be selected once for a given notification. To define a new message domain, click **New ...**
2. Click **OK**. The JMS subscriber will appear in the To list of the Notifications dialog box.

### Iterative address variables

Recipient addresses can also be specified as variables for iterative reporting job step success and failure notifications. In such cases, the variables defined within the step are used to retrieve the recipient address value every time the step is executed.

See the topic “Type for reporting steps” on page 158 for more information.

To add a notification email as an iterative variable, in the Notifications dialog box, click the down arrow next to the To field and then choose **Iterative Subscriber**. The Specify Iterative Variable dialog box appears.

1. Enter the variable in the Iterative Variable field. Use field content assist if necessary. See the topic “Entry field content assist” on page 10 for more information.
2. After all recipients have been selected, click **OK**. The specified variables appear in the To list of the Notifications dialog box.

### Notification message attachments

Content notification messages can include file attachments. If an attachment is specified for a file or a job notification, the new version of the file or the job is included in the message. If an attachment is specified for a folder notification, the object that triggers folder notification (new file or job, or the new version of an existing file or job) is included in the message. Attachments are specified for individual notification recipients separately.

To include a file attachment in a notification message:

1. In the To list, click in the Attachment column next to the notification recipient and choose **Yes** or **No** from the drop-down.

### Customizing the notification message

Notification message can be customized for individual notifications. The subject line and body text of the notification email message are defined by the default message template assigned to the corresponding event type. The Subject and Message fields of the Notification screen contain the default text, variable property values, and for HTML templates, formatting markup. If multiple objects are selected with varying sender, subject, and body text field values, the fields are displayed as blank. Modifying the values of these fields will apply the changes to all selected objects.

1. To change the subject or message body text or formatting, type the changes in the corresponding fields.

*If you make changes to the HTML formatting, you must make sure that the changes are valid. Invalid template may cause notification to fail.*

2. You can use content assist to insert system property variable values into the From field, message subject, or body. See the topic “Entry field content assist” on page 10 for more information.
3. Click **Preview** to display the message. If the template cannot be parsed, an error message is displayed. See the topic “Notification message preview” on page 154 for more information.

The customized message is saved for a specific notification. Customizing messages for individual notifications does not alter the default templates.

*Note:* The appearance and content of the notification message can also be customized by modifying the default template associated the notification event. Contact your IBM SPSS Collaboration and Deployment Services administrator to request global template changes.

## Notification message preview

To preview notification on the notification screen, click **Preview**. The Notification Preview dialog box appears.

The preview displays the email message resulting from merging the template with sample values. The preview includes the specified sender and recipient addresses and customizations made to the message subject, text, and formatting. See the topic “Customizing the notification message” on page 153 for more information. If the customized message cannot be previewed due to invalid formatting, an error message appears.

---

## Subscriptions

IBM SPSS Collaboration and Deployment Services Repository enables the user to subscribe to IBM SPSS Collaboration and Deployment Services files and jobs. Unlike notifications, subscriptions are defined by individual users for themselves.

A subscriber receives email messages when changes occur to the contents of the object and a new version is created. The email message contains a link to the URL of the changed repository object or a file attachment. The Content Explorer also provides the ability to manage subscriptions to individual objects. This may be necessary, for example, when an administrator must remove a subscription for a user who is no longer employed with the company. A user must be assigned corresponding actions through roles to subscribe to content and manage subscriptions.

## Subscribing to files

You can subscribe to a file and specify whether a link to the file URL or the file attachment will be included in the notification message when the file is updated.

To create a subscription to a file, right-click the file name and choose **Subscriptions**. The Create Subscription for Selected File dialog box appears.

1. If the default email address has not been set for the user, click **Update**. The Set User Preferences dialog box appears. Update the address. See the topic “Subscription email address” on page 155 for more information.
2. Specify whether a file link or attachment is included in the email message.
3. Click **OK**.

## Modifying file subscriptions and unsubscribing from files

Use the Modify Subscription for Selected File dialog to modify the subscription settings or to unsubscribe.

To modify a subscription for a file or a IBM SPSS Collaboration and Deployment Services Repository job, or to unsubscribe, right-click the file name and choose **Subscriptions**. The Modify Subscription for Selected File dialog box appears.

**Email.** The email address of the subscriber.

**Link.** Include a link to the URL of the file in the notification message.

**Attachment.** Attach the file to the message.

Modify subscription settings as necessary and click **OK**, or click **Unsubscribe** to remove the subscription.



## Subscription email address

Use the Set email Address dialog to change the default email address for subscriptions.

To change the default email address for subscriptions, in the Create/Modify Subscriptions for Selected File dialog box, click **Update**. The Set Email Address dialog box appears.

**Email Address.** The email address of the user. The address string must contain a @ character and a period. The default address preference for the user will be modified and all existing subscriptions will use the new address. See the topic “Subscriptions recipient” on page 36 for more information.

*Note:* The **Use email address from directory** option is available only for users authenticated by Active Directory and Active Directory with Local Override. The email address from the directory is used by default.

1. Enter the email address or select the **Use email address from directory** option.
2. Click **OK**. The address will appear in the To list of the Notifications dialog box.

## Managing subscriptions

To manage subscriptions to a repository object, right-click the object and choose **Manage Subscriptions**. The Subscriptions for Selected Resource dialog box appears.

The list of subscriptions displays the name of the user and the subscription delivery type (link or attachment). Existing subscriptions can be removed.

1. Select the subscription. Use the **Ctrl** and **Shift** keys to select multiple entries.
2. Click the X button.
3. Click **OK**.

---

## Delivery failures

### Notification Message Delivery Failures

If a notifications email messages cannot be delivered to a specified recipient, a warning message similar to the following will be generated.

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

Subject: IBM SPSS Collaboration and Deployment Services: Job ChurnAnalysis failed  
Sent: 4/5/2010 2:21 PM

The following recipient(s) could not be reached:

jsmith@mycompany.com on 4/5/2010 2:21 PM

The email account does not exist at the organization this message was sent to.  
Check the email address, or contact the recipient directly to find out the correct address.  
<smtp.mycompany.com #5.1.1>

The message is sent to the address indicated in the **From** field of the Notifications dialog box. If the address is invalid, the message is sent to the default IBM SPSS Collaboration and Deployment Services administrator's address. To correct the problem with the notifications delivery, verify that notifications recipients are specified correctly. Notification delivery problems can also occur due to email server setup, network configuration, and so on. If notifications delivery problems persist, contact your system administrator.

### Subscription Message Delivery Failures

If you do not receive messages from subscriptions that have been defined, verify that your default subscriptions address is specified correctly. See the topic “Subscriptions recipient” on page 36 for more

information. If subscription message delivery problems persist, contact your IBM SPSS Collaboration and Deployment Services administrator. The administrator is also notified of any delivery failures by system-generated messages.

---

## Chapter 14. Reporting job steps

A reporting step submits report syntax to an execution server.

For example, a reporting step using an \*.rptdesign report submits **BIRT Report Designer for IBM SPSS** report syntax to the BIRT execution engine. Only the BIRT version shipped with IBM SPSS Collaboration and Deployment Services is supported, and can be installed from the IBM SPSS Collaboration and Deployment Services media or from the IBM SPSS Collaboration and Deployment Services Repository.

To add a reporting step to a job, drag a report from the Content Explorer to the job canvas. However, before the step can successfully run, several properties of the reporting step must be defined. These properties can be classified into several categories:

- General properties, such as the job step name and version of the report.
- Data source properties, including the data source used by the report and its credentials.
- Type properties, such as whether the report type is a single run or an iterative run.
- Parameter properties, such as prompt variable values.
- Result properties, such as output file format and location.
- Clean-up properties, for automatically moving, deleting, or expiring existing output when a job runs.
- Notification properties, for setting up email notifications.

---

### General properties for reporting steps

The general properties for a reporting step specify what runs and which server processes the step when the job runs. To define general properties for a reporting step, click the step in an open job. Click the **General** tab to view or edit the step properties.

**Job Step Name.** Type a name for the step. The default name appends a *\_step* suffix to the file name.

**Report Definition.** Displays the location of the report used for the step.

**Version Label.** Select the version of the report submitted to the execution server using the drop-down list.

---

### Data source for reporting steps

To schedule and run reporting job steps, you need to specify the data source login credentials for the report used in the reporting step. To define data source login credentials, click the step in an open job. Click the **Datasource** tab to view or edit the data source credentials.

Click the Login cell for a data source to select the credentials from those defined in the system.

If the data source for the report uses single sign-on credentials for accessing the data, the report cannot be included in a scheduled job. The credentials for the current user are not stored with the job, and when the job would execute, valid credentials would be needed. The report must specify credentials as a credential defined in the IBM SPSS Collaboration and Deployment Services Repository or as a username/password pair.

In Active Directory or OpenLDAP-based single sign-on environment, Server Process Credential can be used instead of regular user credential. See the topic “Server Process Credential” on page 48 for more information.

If a data source for the report is based on the IBM SPSS Collaboration and Deployment Services Enterprise View, the data provider definition must be specified. Click the Properties cell for the data source to select the data provider definition from those defined in the system. See the topic “Data Provider Definition” on page 217 for more information. Note that you must provide valid repository user credentials for IBM SPSS Collaboration and Deployment Services Enterprise View data source access.

To verify that the specified login information can successfully connect to the data source, click the **Validate Datasource(s)** button. The system reports any problematic login specifications that must be corrected for the job step to run successfully.

---

## Type for reporting steps

To define report type properties, click a reporting step in an open job. Click the **Type** tab to view or edit the step properties.

On the **Type** tab, select the report type for the reporting step. The types available depend on which reporting applications you have installed and which file type you are working with.

### Single

A *single* report runs once and returns a single output file. If the report contains variables, you can specify values for the prompt variables on the **Parameters** tab (see “Parameters for reporting steps” on page 159 for more information). At run time, the report generates output based on the specified variable values.

### Iterative

An *iterative* job runs once but runs in multiple iterations using multiple input parameter mappings. An Iterative Producer report file produces the set of input parameters used for each report iteration. Each iteration produces unique output. For example, to schedule a report to run for each of five departments in your company, you do not need to schedule it five separate times. Instead, you schedule it once and specify variables that run the report multiple times—once for each variable value. The report is scheduled once, but when it runs at the scheduled time, it generates multiple output files—one for each of your five departments.

An iterative job contains an Iterative Producer step and an Iterative Consumer step:

1. Add a new reporting step to the job. The report should contain the list of values to iterate on (for example, a list of departments in your company).
2. On the **Type** tab, select **Iterative Producer**. Type a name for the iterative producer (to be used by the iterative consumer) and select the data set from the report. Verify the report runs successfully.
3. Add a second reporting step to the job.
4. On the **Type** tab, select **Iterative Consumer**. Select the iterative producer name created in step 2. When the job runs, the iterative consumer job step pulls values from the iterative producer job step to produce the report output iterations.

#### Note:

- Iterative Producer job steps cannot contain prompts.
- You must assign Parameters for Iterative Consumer job steps. See “Parameters for reporting steps” on page 159.
- Iterative Producer and Iterative Consumer job steps must use the same data source.
- For BIRT reports, only JDBC data sources are supported for the Iterative Producer—not other BIRT data source types (such as XML).
- BIRT reports can generate iteration source data, or be used to render a report, but not both.

---

## Parameters for reporting steps

If the report contains variables (also known as prompts), you can specify values for the variables on the **Parameters** tab. At run time, the report generates output based on the specified prompt variable values.

The number of prompt variables available for selection is dependent on the number of prompts in the source file and the options used to define those prompts when the file was created. You can only assign prompt variables if your source file contains prompts.

To define parameters, click a reporting step in an open job. Click the **Parameters** tab to view or edit the step properties.

With your cursor in the Assignment column, click the button on the edge of the column. The Prompt Variable Values dialog box will be displayed. Select the variable value to be used for the report. The Required column indicates whether or not the report requires a value for the parameter.

### Note:

- Parameters must be defined for Iterative Consumer job steps. See “Type for reporting steps” on page 158.
- BIRT reports with cascading prompt variables can run in IBM SPSS Collaboration and Deployment Services Deployment Portal but are not supported in IBM SPSS Collaboration and Deployment Services Deployment Manager.
- For BIRT reports, the prompt variable sources must match the data source of the report.

---

## Results for reporting steps

Execution of a reporting step yields report output. To define properties for this output, click a reporting step in an open job. Click the **Results** tab to view or edit the step properties.

The **Format** drop-down identifies the different result types for reporting steps. Properties for each result type include:

**File Name.** Define the file name for the results. You can insert variables into the file output name, and various other fields. To insert variables, position your cursor in the target location in the *File Name* field, type \$, and select one or more values from the drop-down. These variables can help users distinguish between result output. You can insert variables for information such as the date or time. When the report runs as scheduled, the variable information is inserted into the output file name(s). For example, if you have a scheduled report running daily, a user may have trouble determining the information they are viewing. Adding a date variable will insert the current date into the file name at run time, indicating the date the report ran. For more information, see “Entry field content assist” on page 10.

**Location.** Define the location in which to save the results by typing a path or clicking **Browse**. If the specified folder does not exist, it will be created automatically at run time. To insert variables to append the date or time to the folder name at run time, see “Entry field content assist” on page 10.

**Format.** Select the wanted file format for reporting step output. For example, select PDF to produce Portable Document Format (\*.pdf) output. For reports containing images, use the *HTML Archive* format to generate a single archive file containing an HTML report and its ancillary files. This archive can be viewed in a variety of IBM SPSS Collaboration and Deployment Services clients, including IBM SPSS Collaboration and Deployment Services Deployment Portal and IBM SPSS Collaboration and Deployment Services Deployment Manager.

If the output format is specified as HTML for a report with images, an .mht file is generated. This allows all the images and HTML to be combined into a single file. Internet Explorer is able to open and render .mht files automatically.

**Permissions.** Define access permissions for the results by clicking **Browse**. The Output Permissions dialog box will appear. See the topic “Modifying permissions” on page 28 for more information.

**Metadata.** Define output properties for the results by clicking **Browse**. The Output Properties dialog box will appear. See “Output file metadata” on page 128.

**Locale.** Define the locale to use for the results by selecting from the list.

**Rendering Options.** Select from any of the following rendering options.

- Select **Save column headings with data** to display the source column headings with your output.
- Select **Apply edit masks to data** to apply edit masks to your output data. An edit mask is a predefined format users can apply to data in a viewer.
- Select **Output report details only** to display only default data in your output. This option suppresses break group and summary information such as subtotals, grand totals, and counts that are part of your source report.
- Select **Include summary text with data** to display summary text such as subtotals, grand totals, and counts in your output data. Summary text is the descriptive text that displays with the summary line data within a break group. This option is disabled if the **Output report details only** option is selected.

---

## Clean-up for reporting steps

Every execution of a job step produces new versions of any output files saved to the repository. As the number of executions increases, the number of output file versions can quickly become unwieldy, with some versions being obsolete. Clean-up options offer the ability to automatically manage the output of job step executions to control the number of file versions.

Clean-up involves flagging output with automated processing instructions that will be applied to that output the next time the job step executes. When a step executes, the system examines the output rules in place for the output from the last run and modifies the output version as needed when adding a new version from the current execution. For example, suppose a job step is defined to create the file `results.htmlc` in the repository. The clean-up options indicate that the output version should be deleted on the next execution. The first execution of the job step creates the initial version of `results.htmlc` flagged for clean-up. The next execution of the step deletes the last version of `results.htmlc` when creating the new one. Subsequent executions will delete the previous version, allowing only the latest version to appear in the repository.

To define options for cleaning up existing output when a job runs, select a job step in an open job. Select the **Clean-up** tab to view or edit the step's clean-up options. We recommend specifying clean-up options every time you create a new job step, and then disabling or enabling clean-up later as necessary.

Choose from the following options.

- **No clean-up.** The output from the current execution of the step is not flagged with any clean-up options. Subsequent executions of the step will not affect the output from this execution in any way. To specify clean-up options for the output, clear this option.
- **Clean-up Options**
  - **Delete.** The version of any output files from the current job execution will be deleted on the next job execution. The version will be deleted even if the files have been moved to a different folder in the repository since the last run.
  - **Expire.** The version of any output files from the current job execution will be expired on the next job execution. See the topic “Working with Expiration Dates and Expired Files” on page 31 for more information.
  - **Move to.** Output files from the current job execution will be moved to a different folder on the next job execution. Click **Browse** or type a target folder for the moved files. If a file with the same name already exists in the **Move to** folder, it will be overwritten. If the specified folder does not exist, it

will be created automatically during the clean-up process. To use variables to append the date or time to the file name at run time, see “Entry field content assist” on page 10.

- **Label applied to cleaned-up versions**
  - **No label.** No label will be applied to the output version affected by clean-up.
  - **Specify label.** A new specified label will be applied to the output version affected by clean-up.
  - **Select label.** An existing label will be applied to the output version affected by clean-up.
- **On Clean-up Error**
  - **Stop clean-up, do not run the job, mark as failed.** Select this option if you want the job to fail immediately if it encounters errors while attempting clean-up.
  - **Continue clean-up and run the job.** Select this option if you want the job to keep running if it encounters errors while attempting clean-up. If the job runs successfully, the clean-up error will not trigger a failed job history message.
  - **Continue clean-up, run the job, mark as failed.** Select this option if you want the job to keep running if it encounters errors while attempting clean-up. If the job runs successfully, the clean-up error will trigger a failed job history message.

**Important:** Clean-up options are only applied if saving to the repository. In addition, only artifacts flagged for clean-up from the last (most recent) run of the job are affected by clean-up. For example, if an output file has four versions before being flagged for cleanup and the fifth is flagged for clean-up, the fifth version will be handled accordingly during the next execution. The first four versions must be managed manually as no automatic clean-up options were specified during their creation.

---

## Notifications for reporting steps

To define email notification properties, click a reporting step in an open job. Click the **Notifications** tab to view or edit the step properties.

For details about notifications, see “Job step notifications” on page 148.





---

## Chapter 15. Visualization report job steps

A visualization report generates a visual representation of the information contained in a data source. The report is typically created using the IBM SPSS Visualization Designer, which provides an advanced visualization environment enabling users to create graphs ranging from basic business charts to rich, interactive representations using data from a variety of sources. Within the IBM SPSS Visualization Designer, users can access and explore data and define both the structure and the style of graphs. The application provides a range of deployment options for the visualization specifications and resulting graphs; for example, they can be stored in IBM SPSS Collaboration and Deployment Services Repository. Once in the repository, graph specifications can be associated with data as needed, rendered on demand, rendered on a recurring basis, and served to applications and web pages as required.

To add a visualization report step to a job, drag a .viz file from the Content Explorer to the job canvas. However, before the step can successfully run, several properties of the visualization reporting step must be defined. These properties can be classified into several categories:

- General properties, such as the job step name and version of the report.
- Data source properties, including the data source used by the report and its credentials.
- Type properties, such as whether the report type is a single run or an iterative run.
- Parameter properties, such as prompt variable values.
- Result properties, such as output file format and location.
- Clean-up properties, for automatically moving, deleting, or expiring existing output when a job runs.
- Notification properties, for setting up email notifications.

---

### General properties for visualization report steps

The general properties for a visualization report step identify what runs when the job executes. To define general properties for the step, click the step in an open job. Click the **General** tab to view or edit the step properties.

**Job Step Name.** Type a name for the step. The default name appends a *\_step* suffix to the file name.

**Report Definition.** Displays the location of the report used for the step.

**Version Label.** Select the version of the report submitted for processing using the drop-down list.

---

### Data source for visualization report steps

Visualization reports support a variety of input data sources, including:

- Delimited text files
- IBM SPSS Statistics data files (.sav)
- Dimensions Data Models
- IBM SPSS Collaboration and Deployment Services reports
- JDBC sources

Data sources based on files can access files stored on file systems or in the IBM SPSS Collaboration and Deployment Services Repository. For data sources based on reports or JDBC, you need to specify the data source login credentials to access the data. To define data source login credentials, click the step in an open job. Click the **Datasource** tab to view or edit the data source credentials.

---

## Type for visualization report steps

To define report type properties, click a visualization report step in an open job. Click the **Type** tab to view or edit the step properties.

On the **Type** tab, select the report type for the reporting step.

**Single.** Runs once and returns a single output file. If the report contains variables, you can specify values for the prompt variables on the **Parameters** tab (see “Parameters for visualization report steps” for more information). At run time, the report generates output based on the specified variable values.

**Iterative Consumer.** Runs once but runs in multiple iterations using multiple input parameter mappings. An iterative producer report file produces the set of input parameters used by an iterative consumer report iteration. Each iteration produces unique output. For example, to schedule a report to run for each of five departments in your company, you do not need to schedule it five separate times. Instead, you schedule it once and specify variables that run the report multiple times—once for each variable value. The report is scheduled once, but when it runs at the scheduled time, it generates multiple output files—one for each of your five departments.

## Creating an Iterative Visualization Job

An iterative job contains an Iterative Producer step and an Iterative Consumer step:

1. Add a new reporting step to the job. The report should contain the list of values to iterate on (for example, a list of departments in your company).
2. On the **Type** tab, select **Iterative Producer**. Type a name for the iterative producer (to be used by the iterative consumer). Verify the report runs successfully.
3. Add a visualization reporting step to the job, connecting it as a consequent step of the iterative producer using a relationship connector.
4. On the **Type** tab for the visualization report step, select **Iterative Consumer**. Select the iterative producer name created in step 2. When the job runs, the iterative consumer job step pulls values from the iterative producer job step to produce the report output iterations.

---

## Parameters for visualization report steps

Report data sources may contain variables, or prompts, that need an assigned value for proper processing. At run time, the report generates output based on the specified prompt variable values. The number of prompt variables available for selection is dependent on the number of prompts in the source file and the options used to define those prompts when the file was created. You can only assign prompt variables if your source file contains prompts.

To define prompt variables, click a reporting step in an open job. Click the **Parameters** tab to view or edit the step properties.

With your cursor in the Assignment column, click the button on the edge of the column. The Prompt Variable Values dialog box will be displayed. Select the variable value to be used for the report. Select **N** if no value is required for this prompt variable at run time, and you don't want to restrict the report results based on this prompt value (often referred to as a "null" value).

---

## Results for visualization report steps

Execution of a reporting step yields report output. To define properties for this output, click a reporting step in an open job. Click the **Results** tab to view or edit the step properties.

**File Name.** Define the file name for the results. You can insert variables into the file output name, and various other fields. To insert variables, position your cursor in the target location in the *File Name* field,

type \$, and select one or more values from the drop-down. These variables can help users distinguish between result output. You can insert variables for information such as the date or time. When the report runs as scheduled, the variable information is inserted into the output file name(s). For example, if you have a scheduled report running daily, a user may have trouble determining the information they are viewing. Adding a date variable will insert the current date into the file name at run time, indicating the date the report ran. For more information, see “Entry field content assist” on page 10.

**Location.** Define the location in which to save the results by typing a path or clicking **Browse**. The Results Location dialog box will appear. If the specified folder does not exist, it will be created automatically at run time. To insert variables to append the date or time to the folder name at run time, see “Entry field content assist” on page 10.

**Format.** Select the file format for reporting step output. For example, select PDF to produce Portable Document Format (\*.pdf) output.

**Permissions.** Define access permissions for the results by clicking **Browse**. The Output Permissions dialog box will appear. See the topic “Modifying permissions” on page 28 for more information.

**Metadata.** Define output properties for the results by clicking **Browse**. The Output Properties dialog box will appear. See “Output file metadata” on page 128.

**Dimensions.** Define the size of the resulting visualization output by specifying values for the output height and width. Alternatively, select **Use Default Dimensions** to accept the default output size.

---

## Clean-up for visualization report steps

To define options for cleaning up existing output when a job runs, select a job step in an open job. Click the **Clean-up** tab to view or edit the step’s clean-up options.

You should specify clean-up options every time you create a new job step, and then disable or enable clean-up later as necessary.

**Important:** Clean-up options are only applied if saving to the repository.

Choose from the following options.

- **No clean-up.** Output files from the previous run of the current job remain intact.
- **Clean-up Options**
  - **Delete.** Output files from the previous run of the current job will be deleted (even if they have since been moved to a different folder in the repository).
  - **Expire.** Output files from the previous run of the current job will be expired. See the topic “Working with Expiration Dates and Expired Files” on page 31 for more information.
  - **Move to.** Output files from the previous run of the current job will be moved to a different folder. Click **Browse** or type a target folder for old output files. If a file with the same name already exists in the **Move to** folder, it will be overwritten. If the specified folder does not exist, it will be created automatically during the clean-up process. To insert variables to append the date or time to the file name at run time, see “Entry field content assist” on page 10.
- **Label applied to cleaned-up versions**
  - **No label.** Select this option to specify no label is applied to the moved output.
  - **Specify label.** Select this option to specify a new label to apply to the moved output.
  - **Select label.** Select this option to select an existing label to apply to the moved output.
- **On Clean-up Error**
  - **Stop clean-up, do not run the job, mark as failed.** Select this option if you want the job to fail immediately if it encounters errors while attempting clean-up.

- **Continue clean-up and run the job.** Select this option if you want the job to keep running if it encounters errors while attempting clean-up. If the job runs successfully, the clean-up error will not trigger a failed job history message.
- **Continue clean-up, run the job, mark as failed.** Select this option if you want the job to keep running if it encounters errors while attempting clean-up. If the job runs successfully, the clean-up error will trigger a failed job history message.

---

## Notifications for visualization report steps

To define email notification properties, click a reporting step in an open job. Click the **Notifications** tab to view or edit the step properties.

For details about notifications, see “Job step notifications” on page 148.

---

## Chapter 16. SAS® job steps

A SAS step submits SAS syntax to a SAS execution server. This server corresponds to running SAS from the command line, using invocation options to control processing.

Step output consists of a log file detailing the processing history and a text listing file containing the procedure results. The Output Delivery System (ODS) can be used within the syntax to generate HTML results.

To add a SAS step to a job, drag a SAS syntax file from the Content Explorer to the job canvas. However, before the step can be successfully run, several properties of the SAS step must be defined. These properties can be classified into two categories:

- General properties, such as the execution server to process the step
- Result properties, such as output format and location

---

### General properties for SAS steps

The general properties for a SAS step specify what gets run and which server processes the step when the job is run. To define general properties for a SAS step, click the step in an open job. Click the **General** tab to view or edit the step properties.

**Job step name.** Type a name for the step. The default name appends a *\_step* suffix to the file name.

**Object version.** Select the version of the file submitted to the execution server using the drop-down list.

**Additional arguments.** Define optional system options that get passed to the SAS execution server when the job is run.

**Warning expression.** Define warnings for job steps connected by a Conditional connector. The warning expression (e.g., `completion_code`, `warning`, or `success`) must be lowercase.

To use warning expressions:

1. Connect two job steps with a Conditional connector. In the **Expression** field for the conditional connector, type `warning==true`.
2. Navigate to the General tab of the parent job step.
3. In the **Warning expression** field, specify a warning code--for example, `completion_code==18`. This expression overrides the default warning code, if any.

When the job is run, the system will execute the parent job step. Then the system will evaluate the condition for `warning==true`. If true, the system will look at the warning expression specified and determine if the condition was met. If the condition specified in the warning expression was met, the system proceeds to the next job step.

For SAS job steps, the default completion code is `completion_code==1`.

*Note:* If a SAS job step fails because SAS execution server is not defined correctly and the default warning code of 1 is used, the job step status will still be displayed as *Success* in Job History table. This is due to the conflict between Windows warning code 1 (an executable path cannot be found) and SAS default warning code value.

**SAS server.** Use the drop-down list to select a SAS server to process the step. The list contains all servers currently configured to execute SAS steps. To add a new server to the list, click **New**. See the topic “Creating new Content Server connections” on page 14 for more information.

**Remote Process Server.** If SAS is not installed on the machine hosting IBM SPSS Collaboration and Deployment Services Repository, use the drop-down list to select a IBM SPSS Collaboration and Deployment Services Remote Process Server to process the step. The list contains all servers currently configured as remote process servers. The machine hosting SAS needs to be configured as a remote process server and defined as such in IBM SPSS Collaboration and Deployment Services. To add a new server to the list, click **New**. See the topic “Creating new Content Server connections” on page 14 for more information.

## Additional arguments for SAS steps

The additional arguments list for SAS job steps corresponds to the system options that can be supplied to the executable file when invoking SAS.

By default, the additional arguments list includes the two options shown in the following table.

Table 14. Default arguments.

Option	Description
-NOSPLASH	Suppresses the splash screen when invoking the SAS execution server.
-\$NOSTATUSWIN	Suppresses the status window when invoking the SAS execution server.

Additional options can be added to this list for any SAS job step. Each option must be preceded by a hyphen. See the SAS options documentation at for a comprehensive list of available arguments.

## Controlling graph output

SAS steps that create graphs without the use of ODS to save them with HTML output require the specification of a device driver.

Table 15. DEVICE argument.

Option	Description
<b>-DEVICE devicename</b>	Defines devicename as the SAS/GRAPH® output device driver.

For example,

```
-DEVICE Jpeg
```

defines the JPEG driver for creating graphs. In addition to .jpeg, devices commonly used for exporting graphs include .bmp (bitmaps), .gif (gif files), and .wmf (Windows metafiles). To see a list of devices available for your system, consult the *Devices* catalog for the SAS execution server associated with the step.

---

## Results for SAS steps

Execution of a SAS step yields two general types of results: a processing log and an output listing file. To define properties for these results, click a SAS step in an open job. Click the **Results** tab to view or edit the step properties.

The **Name** column identifies the different result types for SAS steps. Properties for each result type include:

**File name.** Specify a name for the results file by clicking the cell for the target and typing the new name.

**Location.** Define the location in which to save the results by clicking the Location cell for the target. Click the resulting ellipsis button to specify the location. Defining result locations for SAS results is identical to doing so for IBM SPSS Statistics results.

**Permissions.** Define access permissions for the results file by clicking the Permissions cell for the output target. Click the resulting ellipsis button to specify the permissions. See the topic “Modifying permissions” on page 28 for more information.

The contents of the output listing file correspond to the output that would be produced in the Output window by executing the step syntax directly within the SAS environment. Although the syntax may produce additional files, such as exported data files or PMML model files, you cannot define properties for those files using a SAS step. If control over the files is needed, use a General job step to execute the SAS syntax.

## SAS result formats

Both the log and listing output are saved as text files. The listing can be saved as HTML by including ODS statements within the SAS program file.

When using ODS to control output format, if the results are saved to the repository, all output files generated by the step are saved to a single compressed file. The name of the file corresponds to the name specified for the Output Target.

### SAS/Graph® results

If the SAS command file associated with a step does not use *ODS* for saving graphs with HTML output, the file must include the proper syntax for exporting the graph output.

1. Use `FILENAME` to create a *fileref* corresponding to the output directory.
2. Assign the *fileref* to the graphics stream using the `GSFNAME` graphic option.

For example, the syntax

```
FILENAME grout 'c:\graph_output';  
GOPTIONS GSFNAME=grout;
```

uses the `c:\graph_output` directory for the graphics stream. To use the current working directory, specify:

```
FILENAME grout '.';  
GOPTIONS GSFNAME=grout;
```

The format of the graph files depends on the `-DEVICE` option specified as an additional argument for the step. See the topic “Controlling graph output” on page 168 for more information.

To save the step results to the repository, use the current working directory as the *fileref* for the graphics stream. IBM SPSS Collaboration and Deployment Services Deployment Manager stores the graph output within a single compressed file in the repository. To save the graph files individually, use a General job step to process the SAS syntax.

---

## Example SAS step

To illustrate the use of a SAS step in a job, this example uses SAS to compute descriptive output for the data introduced in the IBM SPSS Statistics example.

To enable the SAS step to use the data, we transform it to an IBM SPSS Statistics portable file using the following IBM SPSS Statistics syntax:

```
GET FILE='C:\Program Files\data\tutorial\sample_files\customers_model.sav'.  
EXPORT OUTFILE='C:\data\customers_model.por'.
```

Create an IBM SPSS Statistics command file (`sav2por.sps`) containing this syntax and save it to the Content Repository. Create a new job and drag the file to the job canvas.

Click the step to set its properties. On the General tab, specify the IBM SPSS Statistics execution server to process the step. On the Results tab, click the Location cell for the Output Target, click the resulting ellipsis button, and specify a location of *Discard*.

This step is designed to create a new file; the IBM SPSS Statistics output is not of interest.

If the IBM SPSS Statistics step runs successfully, a file named *customers\_model.por* will be created in the *C:\data* directory on the server. We will use the file as input to a SAS step. Consider the following SAS syntax:

```
FILENAME CUST DISK 'C:\data\customers_model.por';
LIBNAME CUST SPSS;
PROC FREQ DATA=CUST._FIRST_;
TABLES response*gender;
RUN;
```

These commands read the IBM SPSS Statistics portable file into a SAS data view and create a crosstabulation of *response* and *gender*. Create a SAS command file containing this syntax, save it under the name *frequencies.sas*, and save it to the Content Repository. Drag the file to the job canvas containing the IBM SPSS Statistics step.

Execution of the SAS step depends on the successful completion of the IBM SPSS Statistics step. To enforce this dependency, we create a *Pass* relationship from the IBM SPSS Statistics step to the SAS step. Select the Pass tool from the job palette, click the IBM SPSS Statistics step, and click the SAS step. With this relationship in place, the *sav2por.sps* step must succeed for the *frequencies.sas* step to be invoked.

Define properties for the SAS step by clicking it. On the General tab, select the SAS server for processing the step. On the Results tab, click the Location cell for the Output Target to define the Location. Click the ellipsis button in the cell to access the Location dialog box. For this example, we save the results to the *Trees* folder in the repository. Click **OK** to return to the Results tab.

Save the job and click the **Run Job Now** button to execute it. When it completes, press the F5 key to refresh the Content Explorer. The output and log files appear in the *Trees* folder.

Suppose we wanted to include graphs in the output. Consider the modified syntax below:

```
FILENAME CUST DISK 'C:\data\customers_model.por';
LIBNAME CUST SPSS;
PROC FREQ DATA=CUST._FIRST_;
TABLES response*gender;
RUN;
```

```
FILENAME grout '.';
GOPTIONS GSFNAME=grout;
```

```
AXIS1
  label=('Frequency')
  MINOR=NONE
;
AXIS2
  label=('Response')
  MINOR=NONE
;
TITLE;
PROC gchart DATA=CUST._FIRST_;
  vbar3d response /
  GROUP=GENDER
DISCRETE
  RAXIS=AXIS1
  MAXIS=AXIS2
  AUTOREF
  SHAPE=B1ock
  PATTERNID=GROUP
  COUTLINE=BLACK
  FRAME
;
RUN;QUIT;

PROC gchart DATA=CUST._FIRST_;
  pie3d response /
```



```
GROUP=GENDER
NOLEGEND
SLICE=OUTSIDE
PERCENT=ARROW
VALUE=NONE
ACROSS=2
COUTLINE=BLACK
;
RUN;QUIT;
```

These commands add bar and pie charts to the output, exporting them to the working folder as defined by the fileref *grout*. Create a new file named *freq\_chart.sas* containing this new syntax and add it to the Content Repository. Delete the *frequencies.sas* file from the job canvas and add *freq\_chart.sas*. Add a Pass relationship from the IBM SPSS Statistics step to the new SAS step.

The inclusion of graphs requires the addition of a system option to define the format. Click the SAS step to modify the properties and click the General tab.

Add the `-DEVICE JPEG` option to save the graphs in the JPG format. Save and run the job. When the job completes, press the F5 key to refresh the Content Explorer.

The *Trees* folder contains the results of the run in the file *freq\_chart.zip*. The listing file contains the results of the `FREQ` procedure. The JPG files *gchart.jpg* and *gchart1.jpg* contain the two graphs. To save the three files individually, use a General job step to process the SAS syntax.



---

## Chapter 17. General job steps

General job steps provide a method for scheduling generic executable processes using IBM SPSS Collaboration and Deployment Services Deployment Manager. These processes may involve execution files or batch files and are often controlled using command line options or switches.

General job steps can access files stored within the repository. As a result, stored data or model files can serve as input to processes. In addition, General job steps provide control over individual output files. For example, individual graphs can be stored in the repository by these steps.

To add a General job step to a job, select the General job tool from the job palette and click the job canvas. However, before the job can run successfully, several properties of the step must be defined. These properties can be classified into three categories:

- General properties, such as the command to run
- Input file properties, such as the location of files used as input to the command
- Output file properties, such as the location of files generated by the command

---

### General properties for general job steps

The general properties for a General job step specify the job to run and the directory in which to place intermediate files.

To define General properties for a General job step:

1. Click a General job step in an open job.
2. Click the General tab to view or edit the step properties.
  - **Job step name.** Type a name for the step.
  - **Iterative Variable List.** If the job step is connected to an iterative producer job step, select the name of the iterative producer from the drop-down list. See the topic “Type for reporting steps” on page 158 for more information.
  - **Command to run.** Type the command to be executed by the step, including the path to the file. Paths containing spaces must be enclosed in quotes or specified using short names, such as *PROGRA~1* for *Program Files*. For diagnostic or troubleshooting purposes, it may be useful to include standard output or standard error redirection in the command specification. Any redirection syntax must match the operating system that executes the command.
  - **Working Directory.** Specify the working directory the step should use when dealing with files. Select **Automatic** to use a custom internal directory, or select **Directory** to specify a particular directory.
  - **Warning expression.** The warning code expression for the job step. The warning expression (e.g., `completion_code`, `warning`, or `success`) must be lowercase. For example, specifying the warning expression `completion_code==26` instructs the system to treat the return code 26 as a success of the job step execution, although it may be different from the default success code.

**Remote Process Server.** If the execution file providing the functionality invoked by the job step is not installed on the machine hosting IBM SPSS Collaboration and Deployment Services Repository, use the drop-down list to select a remote process server to process the step. The list contains all servers currently configured as remote process servers. The machine hosting the execution file needs to be configured as a remote process server and defined as such in IBM SPSS Collaboration and Deployment Services. To add a new server to the list, click **New**. See the topic “Creating new Content Server connections” on page 14 for more information.

---

## Input files for general job steps

The Input Files tab identifies files used as input to the General job step. The step copies an input file from its source location to a specified target location for subsequent processing when running the command.

To define Input File properties:

1. Click a General job step in an open job.
2. Click the Input Files tab to view or edit the properties for input files.
3. Click the **Add** button to add an input file to the step.
  - **File name.** Type the name for the input file. This name should correspond to the name of the file used when the command is run. By default, a new input file is named *inputFile*.
  - **Source Location.** Specify the location of the file. Define this location by clicking the cell and clicking the resulting ellipsis button. See the topic “Input file source location” for more information.
  - **Source Version.** For input files contained within the repository, specify the version of the file to be copied to the target location.
  - **Target Location.** Specify the location to which the step should copy the file. This value should correspond to the location of the file that is accessed when the command is run. Define this location by clicking the cell and clicking the resulting ellipsis button. See the topic “Input file target location” for more information.

### Input file source location

Click the ellipsis button in a selected Source Location cell to define the location of an input file.

- **Location.** Define the location of the input file as either **Content Repository** or **File System**.
- **Path.** Specify the path to the input file. For input files in the Content Repository, click the **Browse** button to navigate to the path.

### Input file target location

Click the ellipsis button in a selected Target Location cell to define the location to which the input file should be copied.

- **Copy to working directory.** Copies the input file from the source location to the working directory defined for the step. See the topic “General properties for general job steps” on page 173 for more information.
- **Copy to another folder.** Copies the input file to a specified folder.

---

## Output files for general job steps

The Output Files tab identifies files generated by the General job step. When the command run completes successfully, the step copies output files from their intermediate locations to specified target locations.

To define Output Files properties:

1. Click a General job step in an open job.
2. Click the Output Files tab to view or edit the properties for output files.
3. Click the **Add** button to add an output file to the step.
  - **File name.** The name of the output file. This name should correspond to the name of the file generated by running the command. By default, a newly added output file is named *outputFile*. To change the name, click in the **File Name** column and type the new name.
  - **Intermediate Output Location.** Specify the location of the file generated by running the command. Define this location by clicking the cell and clicking the resulting ellipsis button. See the topic “Input file source location” for more information.

- **Target Location.** The location of the file. To define location, open the Results Location dialog box by clicking in the column and then clicking the resulting ellipsis button. See the topic “Input file target location” on page 174 for more information.
- **Permissions.** Access permissions for the file if saved to the repository. To define permissions, open the Output Permissions dialog box by clicking in the Permissions column and then clicking the resulting ellipsis button. See the topic “Modifying permissions” on page 28 for more information.
- **Properties.** The properties (metadata) of the file. To define properties, open the Output Properties dialog box by clicking in the Properties column and then clicking the resulting ellipsis button. See the topic “Output file metadata” on page 128 for more information.

## Output file intermediate location

Click the ellipsis button in a selected Intermediate Output Location cell to define the location of the output file created by running the command.

The command can create an intermediate file in one of two locations:

- **Working directory.** Identifies the working directory defined for the step as the location for the intermediate file. See the topic “General properties for general job steps” on page 173 for more information.
- **Write to another folder.** Identifies a specific folder as the location for the intermediate file.

---

## General job step examples

Applications of the General job step are numerous and varied.

The following examples illustrate some of the possible uses.

### Storing IBM Corp. PMML files

The IBM SPSS Statistics syntax in other examples created a PMML file for the generated tree model, but that file could not be stored in the repository by the IBM SPSS Statistics step. A General job step can execute the IBM SPSS Statistics syntax and save all results to the repository.

To create a General job step for storing:

1. Open a new job.
2. Select the General job tool from the job palette.
3. Click in the job canvas.

### Defining General properties

To define General properties for a General job step:

1. Click the step to open the properties dialog box.
2. On the General tab, type a name of *Tree PMML* for the job step.
3. For **Command to run**, type `STATISTICSB -f tree_model.sps -type text -out tree_model.txt`. The switches instruct STATISTICSB to process the syntax file *tree\_model.sps*, saving the results as text in the file *tree\_model.txt*.
4. Select **Automatic** for the **Working Directory**.

### Defining Input File properties

To define Input File properties for a General job step:

1. Click the Input Files tab to define the input for the step. Running the command uses a single input, the syntax file.

2. Add an input file for the step by clicking **Add**. Change the name of this input file to match the name of the syntax file stored within the repository, *tree\_model.sps*.
3. Click the Source Location cell and click the resulting ellipsis button to define the location of the file. The syntax file is stored within the repository, so specify **Content Repository** as the Location.
4. Click **Browse** to navigate to the *Trees* folder to define the path to the file.
5. Click **OK** to return to the Input Files tab.

*Note:* The syntax file itself references a data file, *customers\_model.sav*, stored within the file system on the server. The step does not need to relocate the file so it does not need to be listed as an input. If the data file were stored in the repository or needed to be moved to another system folder for the syntax to access it, it would be listed as an input file for the step.

6. Click the Target Location cell for the input file and click the resulting ellipsis button to define a location to which the step should copy the file. The switch passed to the execution file does not include a path to the syntax file, so select **Copy to working directory** in the Target Location dialog box.
7. Click **OK** to return to the Input Files tab.

## Defining Output File properties

To define Output File properties for a General job step:

1. Click the Output File tab to define the output for the step. Running the command creates a text output file, but we are currently not interested in that file. Our focus is on the PMML file generated by the syntax.
2. Add an output file for the step by clicking **Add**. Change the name of this output file to match the name of the file created by the syntax, *tree\_model.xml*.
3. Click the Intermediate Output Location cell and click the resulting ellipsis button to define the location of the file. The syntax file creates the PMML file in the *C:\Program Files\SPSS* folder, so select **Write to another folder** in the Intermediate Location dialog box and type the path *C:\temp*.
4. Click **OK** to return to the Output File tab.
5. Click the Target Location cell and click the resulting ellipsis button to define the location for the output file.
6. Select **Save to repository** in the Results Location dialog box and click **Browse** to specify the *Trees* folder.
7. Click **OK** to return to the Output File tab.
8. Save the job and click the **Run Job Now** button to execute the job. IBM SPSS Collaboration and Deployment Services Deployment Manager runs the STATISTICSB command and stores the PMML file in the repository. Press the F5 key to refresh the Content Explorer and view the XML file containing the model.

## Batch scoring using IBM SPSS Statistics PMML files

Once a predictive model has been built and the model specifications have been saved as an XML file, the model can be used to score data.

Consider the following syntax:

```
GET FILE='C:\data\customers_new.sav'.
RECODE Age
  ( MISSING = COPY )
  ( LO THRU 37 =1 )
  ( LO THRU 43 =2 )
  ( LO THRU 49 =3 )
  ( LO THRU HI = 4 ) INTO Age_Group.

IF MISSING(Age) Age_Group = -9.
COMPUTE Log_Amount = ln(Amount).

MODEL HANDLE NAME=tree FILE='C:\models\tree_model.xml'
  /MAP VARIABLES=Income_Grp MODELVARIABLES=Income_Group.
```

```
COMPUTE PredCatTree = ApplyModel(tree,'predict').  
SAVE OUTFILE='C:\scores\scoring_results.sav'.
```

This syntax uses the PMML file *tree\_model.xml* to generate scores for the data in *customers\_new.sav*. The results are saved to *scoring\_results.sav*. Save the syntax into a file named *scoring.sps* and add it to a folder named *Trees* in the Content Repository.

To create a General job step for scoring:

1. Open a new job.
2. Select the General job tool from the job palette
3. Click in the job canvas.

## Defining General properties

To define General properties for a General job step:

1. Click the step to open the properties dialog box.
2. On the General tab, type a name of *Scoring* for the job step.
3. For **Command to run**, type `STATISTICSB -f scoring.sps -type text -out scoring.txt`. The switches instruct STATISTICSB to process the syntax file *scoring.sps*, saving the results as text in the file *scoring.txt*.
4. Select **Automatic** for the **Working Directory**.

## Defining Input File properties

To define Input File properties for a General job step:

1. Click the Input Files tab to define the input for the step. This step requires two input files:
  - Syntax file
  - PMML fileBoth files are stored in the *Trees* folder of the Content Repository. To add the files, click **Add** to add an input file for the syntax. Change the name of the input file to match the name of the syntax file stored within the repository, *scoring.sps*.
2. Click the Source Location cell and click the resulting ellipsis button to define the location of the file. The syntax file is stored within the repository, so specify **Content Repository** as the Location.
3. Click **Browse** to navigate to the *Trees* folder to define the path to the file. Click **OK** to return to the Input Files tab.
4. Click the Target Location cell and click the resulting ellipsis button to define a location to which the step should copy the file. The switch passed to the execution file does not include a path to the syntax file, so select **Copy to working directory**. Click **OK** to return to the Input Files tab.
5. Click **Add** to add an input file for the PMML file. Change the name of the input file to match the name of the file stored within the repository, *tree\_model.xml*.
6. Click the Source Location cell and click the resulting ellipsis button to define the location of the file. The syntax file is stored within the repository, so specify **Content Repository** as the Location.
7. Click **Browse** to navigate to the *Trees* folder to define the path to the file. Click **OK** to return to the Input Files tab.
8. Click the Target Location cell and click the resulting ellipsis button to define a location to which the step should copy the file. The syntax requires the model file to be in the *C:\models* folder, so select **Copy to another folder**. Type the path *C:\models* and click **OK**.
9. Click **OK** to return to the Input Files tab.

## Defining Output File properties

To define Output File properties for a General job step:

1. Click the Output File tab to define the output for the step. Running the command creates a text output file and the syntax creates a data file containing the scores.
2. Add an output file for the step by clicking **Add**. Change the name of this output file to match the name of the output file, *scores.txt*.
3. Click the Intermediate Output Location cell and click the resulting ellipsis button to define the location of the file. The command creates the file in the working directory, so select **Working directory** and click **OK** to return to the Output File tab.
4. Click the Target Location cell and click the resulting ellipsis button to define the location for the output file.
5. Select **Save to repository** and click **Browse** to specify the *Trees* folder. Click **OK** to return to the Output File tab.
6. Running the command creates a text output file and the syntax creates a data file containing the scores. To add the files to the Content Repository, add another output file for the step by clicking **Add**. Change the name of this output file to match the name of the data file containing the scores, *scoring\_results.sav*.
7. Click the Intermediate Output Location cell and click the resulting ellipsis button to define the location of the file. The command creates the file in the C:\scores folder, so select **Write to another folder** and type the path C:\scores. Click **OK** to return to the Output File tab.
8. Click the Target Location cell and click the resulting ellipsis button to define the location for the output file. Select **Save to repository** and click **Browse** to specify the *Trees* folder. Click **OK** to return to the Output File tab.
9. Save the job and click the **Run Job Now** button to execute the job. IBM SPSS Collaboration and Deployment Services Deployment Manager runs the STATISTICSB command and scores the new data, storing the results in the repository. Press the F5 key to refresh the Content Explorer and view the file containing the scores.

## Event-based scheduling

Typically, jobs are scheduled for execution based on a particular time-based recurrence pattern—for example, weekly or monthly. In addition, you can also create event-based (or triggered) steps within a job. In this case, your job would contain an event that must be completed successfully before the next step begins. Essentially, an event-based step is a special type of conditional step.

To use an event-based trigger, you create a job step that monitors the existence or the contents of a file. If the file exists or its contents indicate that a new external event has occurred, the job step executes successfully. The completed job step then passes control to the next step in the job. For example, the first step in your job could be a shell script or a batch file that checks to see if a particular file has been updated. If the file was updated, the step completes successfully. The system then proceeds to the next step in the job.

If you use a file in your event-based step, you may need to know the name of the file and specify the file name as a parameter in the script. For example, you may want to use the event name as the file name.

If the event-based step involves a time component, such as a polling interval, it is important to ensure that the polling interval does not conflict with the scheduling recurrence pattern. For example, if your polling interval is one minute and you have scheduled the job to run every hour, then you need to ensure that the monitoring step is complete before the job is scheduled to run—that is, the event-based step has to either succeed or fail before the next job is executed. In this instance, it would be very useful to set up a timeout value in the monitoring script that is less than the scheduling recurrence pattern—for example, a timeout value of 50 minutes.



## Event-based scheduling process overview

Although the individual parameters may vary, the process for implementing an event-based step consists of the following tasks:

1. Create a script that monitors the contents of a file.
2. Create a new job. See the topic “Creating new jobs” on page 119 for more information.
3. Add a General job step to the job that executes the script. You may also need to specify parameters for the script—for example, file name and polling interval. This is the antecedent step. See the topic Chapter 17, “General job steps,” on page 173 for more information.
4. Add another step to the job that depends on the event-based step. This is the consequent step. See the topic “Adding steps to a job” on page 122 for more information.
5. Connect the antecedent step to the consequent step with a pass connector. See the topic “Pass connector” on page 124 for more information.
6. Schedule the job for execution.

## Sample event-based scheduling script

The following is a sample shell script that monitors the contents of a predefined file. This script:

- Assumes that a separate process writes a timestamp to the event file when the given event is completed
- Uses this timestamp to determine whether or not the event has been processed
- Returns an exit code of 0 by default, if the script completes successfully

## Script arguments

The script takes the following command line arguments, which you can specify in the IBM SPSS Collaboration and Deployment Services Deployment Manager when you create the General job step:

- **CONTROL\_FILE**. The script monitors this file for changes. An external process writes events to this file. Periodically, the script checks the file for new events and processes them when found. After a new event has been processed, the script exits successfully. If the script exits successfully, the General job step proceeds to the next step in the job.
- **SLEEP\_TIME**. This argument refers to the amount of time, in seconds, that the script sleeps between checks of the *CONTROL\_FILE*.
- **TIMEOUT\_MINS**. This argument controls the amount of time for which the script should run. If an event has not been processed after the number of minutes specified, script execution stops, and the script returns an exit code of 1. In this example, an exit code of 1 causes the job to fail because the event-based step was not successful. As a result, the job cannot proceed with the remaining steps.

## Script Execution

The script executes the following tasks:

1. The script checks the last row of the *CONTROL\_FILE* file for an event.
2. If the *RESULT\_FILE* exists, the last event has already been processed. The script waits for the designated *SLEEP\_TIME* and then returns to the *CONTROL\_FILE* and repeats the process.
3. If the *RESULT\_FILE* does not exist, then the script creates a new *RESULT\_FILE* for the last event. The script considers the creation of the new *RESULT\_FILE* a success and returns an exit code of 0. The General job step then proceeds to the next step in the job.

```
#!/bin/bash
#-----
#
# Copyright (c) 2005 SPSS, Inc.
# All Rights Reserved
#
# This software is the confidential and proprietary information of
# SPSS, Inc. ("Confidential Information"). You shall not disclose
# such Confidential Information and shall use it only in accordance
# with the terms of the license agreement you entered into with SPSS.
```

```

#
#
#-----

# Control file is written to by an external process.
# This is the file we are monitoring.
CONTROL_FILE=$1

# The duration in seconds we should sleep between
# checks on the control file
SLEEP_TIME=$2

# The timeout (in minutes)
TIMEOUT_MINS=$3

#-----
# usage: Print usage information then exit
#-----
function usage
{
    echo "usage watch.sh <control file> <sleep duration> <timeout in minutes>"
    exit -1
}

#-----
# check_args: Validates the command line arguments
#-----
function check_args
{
    if [ -z $CONTROL_FILE ]; then
        usage;
    elif [ -z $SLEEP_TIME ]; then
        usage;
    elif [ -z $TIMEOUT_MINS ]; then
        usage
    fi
}

#-----
# check_file: Watches $CONTROL_FILE for new events. If a new event occurs,
# touch an event file. If not, sleep for $SLEEP_TIME seconds and try again.
#-----
function check_file
{
    if [ -f $CONTROL_FILE ]; then
        # Get the last event from the file
        LAST_EVENT=`tail -1 $CONTROL_FILE`;
        echo "Last event: $LAST_EVENT"
    else
        "$CONTROL_FILE does not exist. Exiting."
        exit 1
    fi

    # Check to see if we've processed this event
    RESULT_FILE=${0}.LAST_EVENT;
    echo "Checking for event $LAST_EVENT with file $RESULT_FILE";

    if [ -e $RESULT_FILE ]; then
        # We've processed it, sleep and
        # check again later for a new event
        echo "Sleeping for $SLEEP_TIME seconds"
        sleep $SLEEP_TIME;
    else
        # We haven't processed the latest event
        echo "A new event $LAST_EVENT has been fired. Processing.";
        touch $RESULT_FILE;
        FINISHED=finished;
    fi
}

check_args;

# How long we should run before failing
DURATION=${TIMEOUT_MINS*60}

# Get the current time and add our duration to it
NOW=`date +%s`
END_TIME=$((DURATION + $NOW)

```

```

# Finished flag - when set to "finished" we exit
FINISHED=not

echo "Running for a maximum of $DURATION seconds"

while [ $FINISHED != "finished" ] ; do
  check_file;
  NOW=`date +%s`
  if [ $NOW -gt $END_TIME ] ; then
    "Timed out after $DURATION seconds"
    exit 1
  fi
done

# If we've got this far we were successful
# so just return a success code
exit 0

```

## Script failure

Common reasons for a script failure are:

- A new *RESULT\_FILE* is not created before the script times out
- The *CONTROL\_FILE* cannot be found
- Improper specification of command line arguments

If the script fails, it returns an exit code of 1. In this example, this failure will cause the job to fail because the next step in the process cannot be executed.



---

## Chapter 18. Message-based job steps

Message-Based job steps provide a method for triggering processing within an IBM SPSS Collaboration and Deployment Services job with external events signaled by JMS (Java Messaging Service) messages.

Message-based job steps allow to specify message text and message selector as well as the timeout (a time to wait for the message). After the specified message is received, the system proceeds to the next step in the job. If the timeout is exceeded, the job is marked as failed.

To add a message-based step to a job, select the **Message-driven Work** tool from the job palette and click the job canvas. Before the job step can run successfully, its properties must be defined.

---

### General properties of message-based job steps

The general properties for a message-based job step specify the message domain, message text, message selector, and the step timeout.

To define the properties for a message-based job step, click a message-driven step in an open job. Click the General tab to view or edit the step properties.

1. The properties include:

- **Job step name.** Type a name for the step.
- **Message Domain.** The message domain identifies the JMS topic to listen on. See the topic “Message domains” on page 54 for more information.
- **Message Text.** For text message, the text of the expected message of to trigger the job step execution.
- **Message Selector** The selector text to filter messages by contents of the message header.
- **Message Timeout.** The time (in days, hours, minutes, and seconds) to wait for the message triggering the job step. The timeout is effective from the start of the job step after predecessor steps have been completed.



---

## Chapter 19. Notification job steps

Email notifications can be set up in IBM SPSS Collaboration and Deployment Services Deployment Manager using either of the following methods:

- Notifications tab. See the topic “Job success and failure notifications” on page 147 for more information.
- Notification job step.

A notification job step is a specialized method for sending notifications. Typically, the notification job step is used in conjunction with model evaluation.

Notifications can be set up to monitor specific model evaluation outcomes. The following are examples of when notifications might be based on model evaluation return codes:

- Notifications can be set up based on a specific outcome of the model evaluation. For example, if the result returned by a model evaluation analysis is less than 0.85, notifications could be sent to a predefined list of recipients.
- Notifications can be set up based on the traffic light level of a model evaluation job. For example, if the result is red, a notification could be sent. In addition, multiple notification rules could be set up with different email recipients for each traffic light level.

Using a notification job step is similar to specifying notifications as part of another job step. However, the following key differences apply:

- Using notification job steps, different notification templates can be selected for different outcomes. In addition, notification rules can be established for multiple sets of notification recipients.
- In notification job steps, more specific criteria can be established for sending notifications, using conditional connectors. Notifications sent via the Notifications tab apply only for success or failure.
- With notification job steps, all email messages are sent as links. The ability to send an email message as an attachment is not available.

---

### Adding a notification job step to a job

Typically, a notification job step is used in conjunction with another job step.

To add a notification job step to a job:

1. Create a new job or open an existing job.
2. From the job palette, click **Notification**.
3. Click anywhere in the job canvas. The notification job step is added to the job.
4. Because a notification job step is typically connected to another job step, at this stage relationships would be established between the notification job step and other steps within the job. See the topic “Specifying relationships in a job” on page 123 for more information.

For example, if you want to send a notification to a set of recipients if a model evaluation index is less than 80, connect the model evaluation step to a notification step by using a conditional connector. If the evaluation step is successful, the value returned is the index rounded to the nearest integer. Specify the conditional expression for the connector as `success==true && completion_code<80`. If you want to send different emails to different recipients based on index ranges, connect the model evaluation step to multiple notification steps using conditional connectors and specify the index range for each. For example, for three connectors, you could specify the following conditional expressions:

- `success==true && completion_code>90`

- `success==true && completion_code<=90 && completion_code>70`
- `success==true && completion_code<=70`

After the notification job step has been added to the job, its general information and notification recipients must be specified on the corresponding tabs.

---

## General information

The General tab contains information that pertains to the overall notification job step.

The General tab contains the following information:

- **Job step name.** Type a name for the step.
- **Iterative consumer.** Specifies whether or not the notification job step is an iterative consumer of another job step. Select the **Iterative Consumer** check box if this job step is an iterative consumer of another job step. If this check box is selected, one of the available iterative variables must be selected from the drop-down list.

---

## Notification

The Notification tab is used to configure the recipients of notifications for the job step.

The tab contains the following information:

- **Recipients for notifications.** The number of recipients that receive notifications. This number is generated by the system based on the number of recipients specified in the Notifications dialog.

---

## Updating notifications

To update notification parameters:

1. In the General tab, click **Update**. The Notifications dialog box appears.
2. In this dialog box, the following parameters can be modified:
  - **From.** The address of the sender notification message. The field is pre-populated with the default email address in repository configuration options.
  - **To.** The list of notification recipients. The addresses can be entered, specified as an iterative variable for iterative job step notifications, specified as a dynamic list based on a query, or selected from a directory listing of a supported email application, such as Microsoft Outlook or Lotus Notes. To edit an address, click the ellipsis button next to it in the list. Recipients can also be specified as variables. To remove a recipient, highlight the recipient's email address and click **Delete**.
  - **Subject.** The subject of the notification message. By default, the field is populated by the default template associated with the notification event. Modify the message subject, if necessary. See the topic "Customizing the notification message" on page 153 for more information.
  - **Message.** The body text of the notification message. The field is populated by the default template associated with the notification event. Modify the message body text, if necessary. See the topic "Customizing the notification message" on page 153 for more information.
3. Preview the notification message. See the topic "Notification message preview" on page 154 for more information.
4. Click **OK**.

## Selecting a new template

When a notification job step is initially created, the template fields are blank, and a template needs to be selected from the list of available templates. The template selected applies to the current notification job step only and persists until a new template is selected.



In this dialog box, templates can only be modified. New templates cannot be created. To select a new default template:

1. In the Notifications dialog box, click **Select Template**. The Notification Template Selection dialog box appears.
2. Select a template from the list. A preview of the template appears in the Notification Template Selection dialog box. The template can only be viewed in this dialog box. Modifications can only be made in the main Notifications dialog box.
3. Click **OK**. The Notifications dialog box reappears.
4. Make any modifications necessary to the notification message. See the topic “Updating notifications” on page 186 for more information.
5. Click **OK**.



---

## Chapter 20. Champion challenger job steps

---

### Champion challenger overview

By using IBM SPSS Collaboration and Deployment Services Deployment Manager, it is possible to compare model files generated by IBM SPSS Modeler to determine which file contains the most effective predictive model. The champion challenger job step evaluates a model and compares it to one or more challengers.

After the system compares the results, the best model becomes the new champion.

**Champion.** The champion corresponds to the most effective model. For the initial execution of the champion challenger job step, there is no champion--only the first challenger and the corresponding list of challengers. For subsequent executions of the job step, the system determines the champion.

**Challenger.** Challengers are compared against each other. The challenger that generates the best results then becomes the new champion.

### Champion Selection Process

The champion challenger comparison process consists of the following tasks:

1. Scoring each of the competing models.
2. Evaluating the resulting scores.
3. Comparing the evaluation results and determining which challenger is the champion.
4. Saving the new champion to the repository (optional).

### Adding Champion Challenger Work to a Job

To add a champion challenger job step to a job, select the Champion Challenger tool from the job palette and click the job canvas.

### Model evaluation metrics

Model evaluation and comparison can focus on accuracy, gains, or accreditation.

- **Accuracy.** The accuracy of a model reflects the percentage of target responses that are predicted correctly. Models having a high percentage of correct predictions are preferred to those having a low percentage.
- **Gains.** The gains statistic is an indicator of the performance of a model. This measure compares the results from a model to the results obtained without using a model. The improvement in the results when using the model is referred to as the gains. When comparing two models, the model having the higher gains value at a specified percentile is preferred.
- **Accreditation.** Model accreditation reflects the credibility of a model. This approach examines the similarity between new data and the training data on which a model is based. Accreditation values vary from 0 to 1, with high values indicating greater similarity between the predictors in the two data sets. When comparing two models, the model having the higher accreditation value is based on training data that is more similar to the new data, making it more credible and preferred.

### Order dependency

Unlike other types of job steps, the tabs in the champion challenger job step are order dependent.

For example, a challenger must be selected in the Challengers tab before information can be modified in the Champion tab. In addition, the information that appears on some tabs is contingent upon the challengers selected in the Challengers table.

The process for executing a champion challenger comparison consists of the following steps:

1. Providing general job information.
2. Identifying challengers.
3. Specifying champion information.
4. Viewing Enterprise View information and specifying Data Provider Definition assignments.
5. Viewing parameter information.
6. Specifying notifications.

---

## General information

The General tab contains information that pertains to the overall champion challenger job step.

IBM SPSS Modeler server and login information is required to execute a champion challenger job step. Content repository server and login information is required to execute the job and save new champion scenarios to the IBM SPSS Collaboration and Deployment Services Repository. (Content repository server and login information is required, even if you discard the results of the analysis.) Credentials are based upon the user currently logged in to the system.

**Job step name.** The name of the job step. By default, the name of the first job step is *Event 1*. Subsequent job steps are named *Event 2*, *Event 3*, and so on. The name specified here appears in the job history table after the job step is executed.

**IBM SPSS Modeler server.** The IBM SPSS Modeler server or server cluster on which the stream will be executed. The list contains all servers and server clusters currently configured to execute IBM SPSS Modeler steps. To change the server, select from the **IBM SPSS Modeler Server** drop-down list. To create a new server definition, click **New** to launch the server definition wizard.

**IBM SPSS Modeler login.** The credential information used to access the IBM SPSS Modeler server or server cluster. To change the credentials, select a credential definition from the **IBM SPSS Modeler Login** drop-down list. To define new credentials, click **New** to launch the credentials definition wizard.

**Content Repository Server.** The content repository server allows a job to save files to an IBM SPSS Collaboration and Deployment Services Repository. Typically, the content repository server is specified when refreshing models using IBM SPSS Modeler. To specify a content repository server, select a server from the **Content Repository Server** drop-down list. To create a new server definition, click **New** to launch the server definition wizard. To generate a content repository server definition based on the current server information, click **Generate**. A server definition is created and automatically populated in the *Content Repository Server* field.

**Content Repository Login.** The login information for the content repository server. To specify a content repository login, select a credential from the **Content Repository Login** drop-down list. To create a new login, click **New** to launch the content repository login wizard. If single sign-on is not used to connect to the IBM SPSS Collaboration and Deployment Services Repository, click **Generate** to generate a content repository server login based on existing security settings. A content repository login is created and automatically populated in the *Content Repository Login* field. Login generation is not available when using single sign-on.

---

## Challengers

At least one first challenger must be selected to execute a champion challenger job step. It is important to note that the first challenger selected does not imply an order of comparison or any priority in the evaluation process. The first challenger is simply the baseline.

The data source and labels used to determine subsequent challengers are established by the first challenger. After the first challenger is selected or updated, the remaining fields in this tab are updated with information that corresponds to the first challenger.

**First challenger.** Name of the first challenger. To browse through the repository, click **Browse**.

**First challenger label.** Label associated with the model file containing the first challenger. Specify this value when selecting the first challenger.

**Data Source Challenger.** The challenger supplying the data source node used for the job step. Click **Browse** to choose this challenger from the list of entries selected in the Challengers table.

**Metric.** The measurement criteria by which the challenger scenarios are compared. Valid values include *accreditation*, *accuracy* and *gains*. If *gains* is selected, then a percentile needs to be specified as well. See the topic “Model evaluation metrics” on page 189 for more information.

### Challengers Table

The Challengers table lists the default score branches for challengers that match the data source and the label associated with the first challenger. Only the challengers selected from the table will be compared to the first challenger when the job step is executed. Selecting (or clearing) a challenger from the list will cause the system to update the corresponding information on the other job step tabs accordingly.

Each time a job that contains a previously saved champion challenger step is opened, the list of challengers is automatically updated. New challengers may be added to the list if they match the data source and label criteria of the first challenger. Conversely, challengers that no longer meet those criteria may be removed from the list of challengers. If a selected challenger has been removed from the repository, the system will generate a message indicating that the challenger is no longer available.

Although challenger scenarios can be selected and cleared for comparison, information in the challengers table is not modifiable. Specifically, the Challengers table contains the following information.

**Name.** Name of the challenger scenario.

**Label.** Label associated with the challenger scenario.

**Description.** A description of the scenario.

### Modifications to the first challenger

Changes made to the first challenger after the job is saved may affect the champion challenger analysis. For example, suppose that the first challenger is removed from the repository or the label associated with the first challenger is removed. When the Challengers tab is accessed, the system will generate a message, indicating that the first challenger is no longer available for use. In this case, a new first challenger needs to be specified.

### Selecting challengers

To select challenger models for inclusion in champion challenger analyses, perform the following steps:

1. On the Challenger tab of a Champion Challenger step, click **Browse** for the First Challenger. If you are adding challengers manually, click **Add** for the Challengers table.
2. Select the model file by clicking **Browse**. The model file is a IBM SPSS Modeler stream or scenario containing a default score branch with a valid model nugget.
3. Select the label designating the version of the selected model file to use.
4. In the Challengers table, select the score branch to use.
5. Click **OK**.

## Invalid challengers

To be compared, model files must have scoring branches that use a common data structure.

Data characteristics that must match across challengers include the following items:

- Data sources must have the same number of fields.
- The field names must be identical across data sources.
- The field measurement levels must be identical across data sources.

If the system cannot find challengers comparable to the first challenger selected, the Invalid Challenger dialog appears. To select a new challenger:

1. Click **OK** to return to the Challengers tab.
2. Select a new challenger.

## Selecting challenger data sources

To select the data source used in the champion challenger analyses, perform the following steps:

1. On the Challenger tab of a Champion Challenger step, click **Browse** for the data source challenger.
2. From the list of challengers included in the analysis, select the score branch that includes the data source to use.
3. Click **OK**.

---

## Champion

Before information can be specified for a champion, at least one challenger must be selected. If the Champion tab is accessed before a challenger has been selected, the First Challenger Not Selected dialog appears, indicating that a challenger must be selected.

**Do not create a new version of the champion.** Select this option to prevent the creation of a new version of the champion. In this case, the selected labeled version of the champion will be modified. Clear this option to create a new version of the champion instead of modifying the labeled version.

**File name.** The name to use for the copy of the challenger identified as the champion.

**Location.** The location in which the copy of the champion file is stored.

**Permission.** The permissions associated with the copy of the champion.

**Metadata.** Properties associated with the copy of the champion. Specifying metadata for the champion output is the same as specifying metadata for other job output.

## Using the champion in other jobs

After the champion challenger job has been run, the resulting champion can be used in other jobs. To include the champion in another job, the following information is required:

- The name of the champion.
- The location of champion.

When the champion is used in another job, the *LATEST* label is applied. This label cannot be modified.

## Testing the champion

By default, the system creates a new copy of the champion each time the champion challenger job step is run, stores the copy to the output location specified, and writes the results to the job history log.

However, there may be instances in which saving a copy of the champion is not wanted. For example, suppose that you simply want to test the champion challenger job step.

To disable the creation of a copy, select the **Do not create a new version of the champion** check box. If this check box is selected, the remaining options in the tab are disabled. The system will use the same information applied to the current champion.

The system will execute the champion challenger job step and determine a new champion. However, a new version of the champion will not be created or saved to the repository. Instead, the results will only be written to the job history log, indicating which challenger would have been chosen as the champion.

For example, suppose a champion challenger job is run and the creator of the job elected not to create a new version of the champion. The resulting job history log might look like this:

```
Stream execution started
500 500
1000 1000
1500 1500
2000 2000
2500 2500
Stream execution complete, Elapsed=26.22 sec, CPU=18.97 sec
Stream execution started
1000 0
2000 0
Field 'Correct_Sum' has only one value
Field 'Count' has only one value
Field 'Traffic Light Result' has only one value
2855 145
2855 1145
2855 2145
Field 'campaign' has only one value
Field 'gold_card' has only one value
Field 'response' has only one value
Stream execution complete, Elapsed=0.39 sec, CPU=0.2 sec
Stream execution started
500 500
1000 1000
1500 1500
2000 2000
2500 2500
Stream execution complete, Elapsed=26.06 sec, CPU=17.75 sec
Stream execution started
1000 0
2000 0
Field 'Correct_Sum' has only one value
Field 'Count' has only one value
Field 'Traffic Light Result' has only one value
2855 145
2855 1145
2855 2145
```

```

Field 'campaign' has only one value
Field 'gold_card' has only one value
Field 'response' has only one value
Stream execution complete, Elapsed=0.48 sec, CPU=0.19 sec
Stream execution started
500 500
1000 1000
1500 1500
2000 2000
2500 2500
Stream execution complete, Elapsed=21.48 sec, CPU=17.34 sec
Stream execution started
1000 0
2000 0
Field 'Correct_Sum' has only one value
Field 'Count' has only one value
Field 'Traffic Light Result' has only one value
2855 145
2855 1145
2855 2145
Field 'campaign' has only one value
Field 'gold_card' has only one value
Field 'response' has only one value
Stream execution complete, Elapsed=0.39 sec, CPU=0.17 sec
The result for challenger cc_cartresponse.scn is 98.809.
The result for challenger cc_neuralnetresponse.scn is 98.844.
The result for challenger cc_c51response.scn is 98.809.
The declared Champion is cc_neuralnetresponse.scn.

```

Note the last line in the log file:

The declared Champion is `cc_neuralnetresponse.scn`.

This line indicates that the `cc_neuralnetresponse` scenario would have been the champion. However, a copy of this scenario was not saved to the repository because the system did not create a new version of the scenario. If the scenario had been saved to the repository, the log would contain an additional line, indicating that the scenario was saved to the repository--for example:

Adding artifact `spsscr:/PMDemo/ModelManagement/cc_neuralnetresponse.scn`.

---

## Data files

Data files information appears for the challengers that were selected from the list of challengers in the Challengers tab.

Any changes made in this tab apply to the champion challenger job step only. Modifications made to the data file information are not propagated back to the challenger saved in the IBM SPSS Collaboration and Deployment Services Repository. The data files table contains the following information.

**Node Name.** The name of the input node that contains the data used by the stream. The node name is not modifiable.

**Node Type.** The node type as defined in the stream. The node type is not modifiable.

**File name.** The name of the input data file. To change the name, click in the file name cell and change the name.

**Format.** The format of the output file--for example, a comma delimited file. To modify the file format type, click in the Format cell. A drop-down arrow appears. Select the format type.



**Location.** The location of the input data files. To modify the location, click in the column and then click the resulting ellipsis button. The Input File Location dialog box opens. Change the location as necessary.

---

## Data view

Analytic data view information appears for the challengers that were selected from the list of challengers in the Challengers tab.

Any changes made in this tab apply to the champion challenger job step only. Modifications made to the data view information are not propagated back to the challenger saved in the IBM SPSS Collaboration and Deployment Services Repository. The data view table contains the following information.

**Node Name.** The name of the data view node that contains the data used by the stream. The node name is not modifiable.

**Analytic Data View.** The analytic data view referenced by the data view node.

**Label.** Label identifying the version of the analytic data view used.

**Table Name.** The table containing the input data fields.

**Data Access Plan.** The plan supplying the data records for the input data fields. To change the data access plan used for a node, select the cell containing the access plan and click the resulting ellipsis (...) button.

---

## Enterprise view

The Enterprise View tab specifies the Enterprise View settings for accessing data. This tab lists only the node associated with the default branch of the model files.

**Enterprise View Node.** Name of the Enterprise View node in the stream or scenario.

**Data Provider Definition.** Data provider definition associated with the Enterprise View node. To change the data provider used for a node, select the cell containing the data provider and click the resulting ellipsis (...) button.

**Label.** Label identifying the version of the data provider used.

---

## ODBC data sources

ODBC data source information appears for the challengers that were selected from the list of challengers in the Challengers tab.

Any changes made in this tab apply to the champion challenger job step only. Modifications made to the ODBC data source information are not propagated back to the challenger saved in the IBM SPSS Collaboration and Deployment Services Repository. The ODBC data sources table contains the following information.

**Node Name.** The name of the input node that contains the data used by the stream. The name is prefixed by the names of any supernodes containing the node separated by slashes. For example, if the node *MyNode* is within a supernode named *Supernode1*, the name appears as */Supernode1/MyNode*.

**Node Type.** The node type as defined in the stream.

**ODBC Data Sources.** The current ODBC data source name (DSN). To change to a different ODBC data source, click the cell containing current data source name, then click the "..." button that is displayed. Doing so displays a dialog box from where you can choose an existing DSN or create a new one.

**Credentials.** To change the database username and password when changing the ODBC data source, click the cell containing the current credentials, then click the "..." button that is displayed. Doing so displays a dialog box from where you can choose an existing credential definition or create a new one.

**Database Table.** The database table that corresponds to the node.

Nodes within locked supernodes are not accessible. They cannot be viewed or modified.

---

## Cognos import

If the model files contain any IBM Cognos BI source nodes, the Cognos connection details are displayed here.

**Node Name.** The name of the Cognos source node.

**Connection URL.** The URL of the Cognos server to which the connection is made.

**Package Name.** The name of the Cognos package from which the metadata is imported.

**Anonymous.** Contains **Anonymous** if anonymous login is used for the Cognos server connection, or **Credential** if a specific Cognos username and password are used.

**Credentials.** The username and password (if required) on the Cognos server.

*Note:* Cognos credentials must be created in a domain, which represents the Cognos namespace ID.

---

## Chapter 21. Submitted jobs

The Submitted Jobs folder is a staging area in IBM SPSS Collaboration and Deployment Services Repository that displays the results of running reports using IBM SPSS Collaboration and Deployment Services Deployment Portal. From the reports, jobs and additional output are generated and appear in the Submitted Jobs folder.

For example, Business Intelligence and Reporting Tools (BIRT) reports can be submitted using IBM SPSS Collaboration and Deployment Services Deployment Portal, and their associated jobs and output appear in Submitted Jobs. For details about BIRT, see <http://www.eclipse.org/birt>.

The Submitted Jobs folder is at the same level as the Content Repository in the Content Explorer. The Submitted Jobs folder appears in the Content Explorer, whether or not content is available within it.

When a user submits a report to IBM SPSS Collaboration and Deployment Services, a folder is created for the user within the Submitted Jobs folder. The folder name corresponds to the user's login ID. Subsequently, output is placed into the folder. Each report is contained in a separate folder, whose name contains a timestamp. Each report folder contains a corresponding job and any additional artifacts related to the report.

The following guidelines apply to objects in the Submitted Jobs folder:

- The objects within the Submitted Jobs folder are saved to the IBM SPSS Collaboration and Deployment Services Repository.
- The contents of the Submitted Jobs folder are searchable. See the topic “Searching” on page 17 for more information.
- Objects, such as files and jobs, can be opened but only in read-only mode.

For more information about reporting, see Chapter 14, “Reporting job steps,” on page 157.

---

### Restrictions within the Submitted Jobs folder

Objects in the Submitted Jobs folder cannot be modified—for example, names cannot be changed, jobs cannot be scheduled, and versioning is disabled. To modify an object, it must be moved from the Submitted Jobs folder to the Content Repository. Once an object is moved into the Content Repository, full functionality is restored.

Although objects can be moved from the Submitted Jobs folder to the Content Repository, the reverse is not possible. Objects from the Content Repository cannot be moved into the Submitted Jobs folder.

The only exception to these restrictions is permissions. Permissions can be changed for objects within the Submitted Jobs folder. All of the standard permissions are available; however, only the owner of an object can modify its permissions. For example, the owner can grant permissions to modify an object. By default the user who created the report is the owner. See the topic “Modifying permissions” on page 28 for more information.

---

### Submitted jobs and expiration dates

By default, objects within the Submitted Jobs folder automatically expire after a set period of time. To prevent an object from expiring, it must be moved out of the Submitted Jobs folder and into the Content Repository.

When the object is moved into the Content Repository, the expiration date attached to the object is not automatically removed. To prevent the object from expiring, the expiration date must be explicitly removed or, if needed, a new expiration date must be set.

The expiration time period in the Submitted Jobs folder is different from setting a specific expiration date in the Content Repository. Unlike expired objects in the Content Repository, expired objects in the Submitted Jobs folder are automatically deleted from the repository. The expiration time period for the contents of Submitted Jobs folder is a configuration setting, which can be modified by an administrator. The default value is five days. Expiration dates in the Content Repository are user-defined and modifiable. See the topic “Working with Expiration Dates and Expired Files” on page 31 for more information.

---

## Chapter 22. Accessibility features

Accessibility features help a user who has a physical disability, such as a visual impairment or limited mobility, to use software products successfully.

This section provides an overview of alternative methods for accessing the functionality of the product. More specifically, the following topics are covered:

- Keyboard navigation of the user interface
- Special issues for visually impaired users
- Special issues for blind users

---

### Keyboard navigation

Keyboard shortcuts allow you to navigate the user interface without using a mouse. At the most basic level, you can press the Alt key plus the appropriate key to activate window menus or press the Tab key to scroll through dialog box controls.

The following table describes general shortcuts that may be used throughout IBM SPSS Collaboration and Deployment Services Deployment Manager.

*Table 16. General keyboard shortcuts.*

Shortcut	Action
Shift + F10	Brings up a menu, similar to a right-click.
Ctrl + F7	Brings up a list of active views and allows you to select one.
Tab	Iterates through the different controls on the view, such as tables or toolbars
Shift + Tab	Iterates through the different controls on the view in reverse.
Ctrl + Tab	Advances to the next control on the view when focus is on a control that accepts a Tab as input. For example, the Message text box of the notification dialog boxes allows tab entries. To move to the next control instead of inserting a tab character, use Ctrl+Tab.
Ctrl + M	Maximizes the active view or returns a maximized view to its original size
Alt	Highlights mnemonics. The menu titles in the menu bar have their mnemonic letters underlined.
Alt + mnemonic	Selects the control that is mapped to the mnemonic, such as a menu item, button, check box, or text field. For example, Alt + F opens the <i>File</i> menu. Pressing Alt + N in the Roles editor brings up the Create New Role dialog box.

### Navigating the content explorer

The following table describes keyboard shortcuts that may be used within the Content Explorer.

*Table 17. Content Explorer shortcuts*

Shortcut	Action
Up/down Arrows	Highlights items in the tree.
Enter	Opens the selected item.

## Navigating tables

The following table describes shortcuts that may be used within tables appearing in views and dialog boxes. Use the Tab key to select the table for subsequent navigation.

Table 18. Table shortcuts

Shortcut	Action
Up/down Arrows	Highlights items in the tables.
Right arrow	Expands a selected item, if the item can be expanded.
Left arrow	Collapses an expanded item.
Ctrl+<column number>	For a highlighted row, moves focus to the cell in the specified column if the cell value can be modified. Both keys must be released simultaneously. If the value cannot be modified, the focus is unchanged.

## Navigating the job history and job schedule views

The following table describes keyboard shortcuts that may be used within the job history and job schedule views.

Table 19. Job History and Job Schedule shortcuts.

Shortcut	Action
Up/down Arrows	Highlights items in the tables.
Right arrow	Highlights items in tables. Expands a selected item, if the item can be expanded.
Left arrow	Collapses an expanded item.
Enter	Opens a selected item in the editor. (Note: For jobs, use Shift + F10 to open a menu for a selected Job and select <b>Open In Job Editor</b> .)

## Navigating the job editor

The following table describes keyboard shortcuts that may be used within the job editor.

Table 20. Job Editor shortcuts.

Shortcut	Action
Up/down Arrows	Highlights options in the job palette. Moves the cursor around in the graphical editor. If the Move anchor is selected, repositions a job step in the editor.
Left-arrow/right-arrow	Moves the cursor around in the graphical editor. Cycles the selection through the different job steps and the event flow items as well. If the Move anchor is selected, repositions a job step in the editor.
Period (.)	Cycles through the move options for a selected job step.
Plus sign (+)	Adds a new step of the type that is selected in the palette to the job.

## Navigating the help system

The following table describes keyboard shortcuts that may be used within the help system.

Table 21. Help system shortcuts.

Shortcut	Action
Up/down Arrows	On the Contents tab, moves up and down through the topics.
Right Arrow	On the Contents tab, expands any subtopics of the current topic.
Left Arrow	On the Contents tab, moves to the parent of the current topic. If the current topic has expanded subtopics, the subtopics are closed.

Table 21. Help system shortcuts (continued).

Shortcut	Action
Shift+F10	On the Contents tab, displays the menu from which the selected topic can be opened.

---

## Accessibility for visually impaired users

IBM SPSS Collaboration and Deployment Services Deployment Manager honors the contrast settings specified for your operating system. To adjust these settings, consult your operating system documentation.

In addition, you have control over the magnification in PDF files viewed in Acrobat Reader. To set the magnification in Acrobat Reader:

1. From the View menu, choose the **Zoom To** option on the Zoom submenu.
2. Specify the magnification level.

---

## Accessibility issues for blind users

Support for blind users is predominately dependent on the use of a screen reader.

IBM SPSS Collaboration and Deployment Services has been fully tested with a screen reader. For information on installing and configuring your screen reader, consult the vendor documentation.





---

## Chapter 23. IBM SPSS Collaboration and Deployment Services Enterprise View

---

### IBM SPSS Collaboration and Deployment Services Enterprise View overview

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

### IBM SPSS Collaboration and Deployment Services Enterprise View builder interface

The Enterprise View builder user interface is an extension to the IBM SPSS Collaboration and Deployment Services Deployment Manager interface. The Enterprise View builder includes a set of editing and administration interfaces used by data experts to create and manage the Enterprise View.

The Enterprise View builder user interface consists of:

- **Enterprise View editor**—Records information about the entities and attributes (tables and columns) used in predictive applications. In a given repository, the Enterprise View is a singleton object; it may have versions, but there cannot be more than one version tree. See the topic “Enterprise View” on page 204 for more information.
- **Application View editor**—Overlays a subset of the Enterprise View, recording how the Enterprise View tables and columns are used in the context of a specific application (for example, call center or fraud detection). See the topic “Application View” on page 212 for more information.
- **Data Provider Definition editor**—Records information needed to access business data in a particular context. This allows data to flow from different sources depending on what the using-application does with the data. For example, the values for CustomerID might flow from a warehouse table in an IBM SPSS Modeler stream used to build an analytic, while the values might flow from physical data source. The Data Provider Definition supplies the information needed at runtime to access data. See the topic “Data Provider Definition” on page 217 for more information.
- **Data Provider Definition - Real Time editor**—Used exclusively for real time server interactions with applications. For example, the values for CustomerID might flow from a warehouse table in an IBM SPSS Modeler stream used to build an analytic, while the values might flow from a user-interface (via a data object) when the scorable analytic produced by the stream is used in a real-time scoring scenario. Unlike a Data Provider Definition, a Data Provider Definition - Real Time can only retrieve one row of data at a time. See the topic “Data Provider Definition - Real Time” on page 221 for more information.

## Credential definitions

IBM SPSS Collaboration and Deployment Services Enterprise View objects, namely the Enterprise View and Data Provider Definition, require access to data from physical data sources. In order to log in to these data sources, credential definitions must be created for each data source. Each credential definition contains a user ID and password. See the topic “Adding new credentials” on page 47 for more information.

## Data source definitions

IBM SPSS Collaboration and Deployment Services Enterprise View builder requires both an ODBC and JDBC data source definition for every data source. Both definition types should point to the same data source. See the topic “Data sources” on page 48 for more information.

Keep in mind the following:

- The Data Provider Definition - Real Time requires only JDBC data sources.
- For ODBC data source definition types, the referenced ODBC data source names (DSNs) must exist on the same system on which the IBM SPSS Collaboration and Deployment Services Repository is installed.
- When configuring DSNs, regardless of type (Oracle, SQL Server, DB2, and so on), ensure that the **quoted identifier** option is enabled (if available).
- When configuring DSNs, regardless of type (Oracle, SQL Server, DB2, and so on), you must provide the appropriate database information.
- Third-party JDBC drivers, such as for Netezza or Teradata, can be used by augmenting your installation with the driver files. See the topic “Third-party JDBC drivers” on page 51 for more information.

---

## Enterprise View

The Enterprise View is an abstracted dictionary of the enterprise data used in predictive analytic applications. It provides a simplified business-oriented picture of the enterprise data available for data mining and reporting. The Enterprise View is constructed and maintained by a data expert (a database administrator, for example). The Enterprise View editor in IBM SPSS Collaboration and Deployment Services Deployment Manager provides the data expert with a graphical user interface to build, maintain, and validate global information about enterprise data. Taken together, the Enterprise View editor, the Application View editor, the Data Provider Definition editor, and the Data Provider Definition - Real Time editor constitute the IBM SPSS Collaboration and Deployment Services Enterprise View builder user interface used by data experts in constructing, validating, and managing the IBM SPSS Collaboration and Deployment Services Enterprise View.

Upon installation, the Enterprise View is automatically created as an empty singleton object alongside the IBM SPSS Collaboration and Deployment Services Repository. The data expert populates the empty Enterprise View with table and column metadata used in predictive analytic applications. Changes to the Enterprise View are handled in the same manner as all repository objects; a new version is created and all previous versions continue to exist, unless they are deleted manually.

It is not possible to delete the Enterprise View from the repository completely. At least one version is required. Only users who are members of the **Administrators** role, and those assigned to a role with the **Manage Enterprise View** action, are allowed to work with the Enterprise View.

## Enterprise View editor

The Enterprise View editor allows data experts to work with the Enterprise View definition. The editor provides options for creating, editing, and removing the logical entities and attributes required by enterprise level applications. Data experts can create table and column entries manually or import table/column information from physical database tables.

*Note:* Only users with **ADMIN** authority, and those assigned to a role with the proper authority, are allowed to work with the Enterprise View.

## Adding tables to the Enterprise View

Clicking **Add Table** in the Enterprise View editor launches the Add Table wizard. The wizard provides options for specifying a new table name and accompanying description, and allows you to create an empty table or copy the table structure from an existing data source.

**Name.** Enter an appropriate table name. Table names are limited to 256 characters and cannot contain leading or trailing spaces. Each Enterprise View table name must be unique.

If you will use a Classic Federation Server as a data source associated with a table, specify the table and column names using uppercase characters to ensure proper matching between the physical and logical definitions.

**Enter a description (Optional).** Enter an appropriate table description.

**Create an empty table.** When this option is selected, you must provide a table name and then click **Finish** to create an empty table.

**Select columns from a physical table.** When this option is selected, additional wizard steps become available, allowing you to define additional table attributes.

1. After entering the appropriate information, click **Next** to continue to the “Adding tables/columns: Connection settings” step.

**Adding tables/columns: Connection settings:** The Connection Settings dialog box provides options for selecting a JDBC data source provider and data source credentials. The dialog also provides options for creating new data sources and credentials. Refer to “Selecting a data source definition type” on page 49 and “Adding new credentials” on page 47 for more information.

**JDBC.** Select an appropriate JDBC data source definition.

**Credentials.** Select an appropriate credential definition.

**Connect.** You must click **Connect** to verify the defined data source connection.

1. If the connection step does not return any errors, click **Next** to continue to the “Adding tables/columns: Selecting columns to include” step.

**Adding tables/columns: Selecting columns to include:** The Select Columns to Include dialog box provides options for selecting an existing data source table and its associated column definitions.

**Table Types...** Click to select what to return from the physical data source.

- **System Table**—Returns data source system tables (for example, sysdevices and sysfiles).
- **Table**—Returns regular data source tables (for example, CUSTOMERS and ORDERS).
- **View**—Returns table views. Views are virtual tables whose contents are defined by a query. As with a real table, a view consists of a set of fields and records. The records and fields come from tables referenced in the query defining the view and are produced dynamically when the view is referenced.

*Note:* Other options may be available, depending on the data source configuration.

**Note:** When using the Add Table wizard, you can select or clear any available physical column. When using the Add Column wizard, you can select and clear only physical columns that are not already defined for the Enterprise View. For example, if the column `CUSTOMER_ID` from physical table `CUSTOMERS` is already defined in the Enterprise View, you cannot select or clear the column in the Select Columns to Include dialog box.

1. After selecting the appropriate table and columns, click **Finish** to return to the Enterprise View editor.

**Note:** By default, all available columns are selected automatically.

The new table appears in the table section (when using the Add Table wizard), and the selected column definitions that were copied from the defined data source are displayed in the column section.

## Removing tables from the Enterprise View

To remove existing tables:

1. Select the table.
2. Click **Remove Table**. A dialog box opens, informing you that the table and all associated columns will be deleted and that the action cannot be undone.

## Adding columns to the Enterprise View

Clicking **Add Column** from the Enterprise View editor launches the Add Column wizard. The wizard provides options for specifying a new column name and accompanying description, and allows you to create a single column or select the column structure from an existing data source.

**Name.** Enter an appropriate column name. Each Enterprise View column name must be unique. Column names are limited to 256 characters and cannot contain leading or trailing spaces.

When creating column names, you should take into account the enterprise data environment. The type of database that will be mapped to the columns can influence the lengths of column names. For example, when an Oracle database supplies the data for the columns, the column names can have a maximum of 30 characters due to restrictions within that database. Column names exceeding this length will need to be modified to accommodate the physical data.

If you will use a Classic Federation Server as a data source associated with a table, specify the table and column names using uppercase characters to ensure proper matching between the physical and logical definitions.

**Type.** Select the appropriate column data type. Each column can store data consisting of only a single data type. See the topic “Type” on page 208 for more information.

**Enter a description (Optional).** Enter an appropriate column description.

**Create a single column.** When this option is selected, you must provide a column name and type and then click **Finish** to create an empty column.

**Select columns from a physical table.** When this option is selected, additional wizard steps become available, allowing you to define which columns to add from the physical data source.

1. To create a column manually, select **Create a single column**, enter an appropriate column name and type, and then click **Finish**. You are returned to the Enterprise View editor and the new column for the selected table.
2. To copy one or more column definitions from a physical data source, select **Select columns from a physical table** and click **Next** to continue to the “Adding tables/columns: Connection settings” on page 205 step.

**Sorting columns:** All columns can be sorted in ascending or descending order. To sort a column, click an appropriate column heading from the column section. The arrow displayed in the heading indicates whether the column is sorted in ascending or descending order.

*Note:* Sorting does not affect how columns are displayed for consuming applications. The index order dictates how columns are displayed for consuming applications. See the topic “Index” on page 213 for more information. .

## Removing columns from the Enterprise View

To remove existing columns:

1. Select the columns.
2. Click **Remove Column**. A dialog box opens, informing you that the column will be deleted and that the action cannot be undone.

## Changing Enterprise View column attributes

Column attributes can be changed by double-clicking the column cell. You can also select multiple rows, right-click, and apply the same attribute changes to all selected rows (for example, you can apply the same **Type** to all selected rows).

**Note:** Changing row attributes via multirow selection erases the undo/redo history.

**Index.** The **Index** column represents the order in which the Enterprise View columns will be presented to consuming applications. See the topic “Index” for more information.

**Column Name.** Displays the current column name. See the topic “Column name” for more information.

**Type.** Specifies the type of data stored in a column. Each column can store data consisting of only a single data type. See the topic “Type” on page 208 for more information.

**Categories.** The categories attribute allows you to define categorical values for a column. See the topic “Categories” on page 208 for more information.

**Minimum/Maximum.** The minimum/maximum attributes function as user aids. See the topic “Minimum/Maximum” on page 209 for more information.

**Description.** Allows you to enter detailed table/column description text. See the topic “Description” on page 209 for more information.

**Index:** The column index order dictates how columns are displayed for consuming applications. In the following table, the columns will be displayed to consuming applications in the order *Income, State, Ownership, Mortgage, and Second Mortgage*.

Index	Column name
4	Mortgage
2	State
5	Second Mortgage
3	Ownership
1	Income

By default, the index order is determined by the order in which the columns were defined in the Enterprise View. The index order can be edited manually at the Enterprise View or the Application View level.

**Note:** When the index order is changed at the Application View level, it does not affect the Enterprise View indexing order.

**Column name:** The column name attribute displays the current column name inherited from the Enterprise View.

When creating column names, you should take into account the enterprise data environment. The type of database that will be mapped to the columns can influence the lengths of column names. For example,

when an Oracle database supplies the data for the columns, the column names can have a maximum of 30 characters due to restrictions within that database. Column names exceeding this length will need to be modified to accommodate the physical data.

**Type:** The type attribute specifies the type of data stored in a column. Each column can store data consisting of only a single data type. See the topic “IBM SPSS Collaboration and Deployment Services Enterprise View data types” on page 254 for more information.

Table 22. Data types

Type	Description
Boolean	Used to store binary digits (0 or 1). These can be used to denote Boolean values, such as Yes/No, True/False, or On/Off.
Date	Values must have the form CCYY-MM-DD, where CCYY denotes the year, MM denotes the number of the month (January = 1, February = 2, and so on), and DD denotes the day (starting with 1). Optionally, the year can be preceded by + or -. Leading zeros for the year, month, and day are optional. Spaces are not allowed before or after the - separator.
Daytime	Values must have the form HH:MM:SS.xxx, where HH denotes the hours (ranging from 0 to 23), MM denotes the minutes (0 to 59), SS denotes the seconds (0 to 59), and xxx, the milliseconds (0 to 999). First, note that only the 24-hour representation is allowed. Furthermore, the seconds and milliseconds are optional. Leading zeros for the hours, minutes, and seconds are optional. Spaces are not allowed before or after the : separator.
Decimal	Used to store exact numbers. The scale and maximum precision can be configured. The precision denotes the maximum number of digits each value may consist of. The scale indicates how many of those digits may occur after the decimal point (that is, the fractional part). The value of this parameter must be greater than or equal to zero and less than or equal to the precision.
Double	Represents signed, approximate numerical values. Double permits much larger numbers than Float.
Float	Represents signed, approximate numerical values. Double permits much larger numbers than Float.
Integer	Used for whole numbers (for example, 5 or 110).
Long	Used for character data of variable length up to two gigabytes.
String	Used to store a sequence of ISO-8859-Latin-1 characters. The maximum length of such a sequence is configurable. Furthermore, the string may be empty and cannot contain the newline (ASCII 10) or null (ASCII 0) characters.
Timestamp	Allows values of the form CCYY-MM-DDTHH:MM:SS.xxx in accordance with the ISO 8601 standard. It is easy to see that it is nothing more than the concatenation of the Date and the Time types, using a delimiter of T to indicate the beginning of the Time portion. Consequently, the rules for the Date and the Time types are also applicable to the Timestamp type.

**Categories:** The categories attribute allows you to define categorical values for a column. The category attributes are used by consuming applications to populate the values for selection fields or menus. Categorical values can be added to any column regardless of its Type.

1. Double click inside a Categories column cell to display the ... icon. Click the icon to display the Categorical Values dialog box.

### Defining Categories

1. Click **Add...** to launch the Add Categorical Values wizard.

**Single Value.** Select this option to manually enter a single categorical value. After you enter an appropriate value, click **Finish** to add the entry to the Categorical Values list.

**Import from file.** Select this option to import categorical values from an existing .csv file.

2. Click ... to display the Select File dialog box.
3. Locate the appropriate .csv file and click **Open**. You are returned to the Add Categorical Values dialog box, and the path for the selected .csv file displays in the **Import from file** field.
4. Click **Finish** to add every unique value in the .csv file to the Categorical Values list.  
**Import from data source.** Select this option to import categorical values from an existing data source table column.
5. Click **Next** to display the Connection Settings step.
6. Select the appropriate data source **JDBC** and **Credentials** entries, click **Connect**, and then click **Next** to display the **Select Table and Column** step.  
*Note:* You may need to create new data source and credential definitions if they do not already exist.
7. Click **Finish** to add the categorical values, for the selected column, to the Categorical Values list.
8. Click **Finish** to return to the Enterprise View editor. The selected column now displays the categorical values you defined.

*Exporting categorical values:* The Categorical Values dialog box allows you to export your defined categorical values to a .csv file. To export your defined categorical values:

1. Click **Export**. The Export dialog box displays.
2. Navigate to an appropriate save location.
3. Enter a suitable file name for the .csv file.
4. Click **Save**.

**Note:** The export process saves all defined categorical values, regardless if they are all selected in the Categorical Values dialog box.

5. Click **OK** to return to the Enterprise View editor.

*Removing categorical values:* To remove existing categorical values:

1. Select the appropriate values from the Categorical Values dialog box.
2. Click **Remove**.
3. Click **OK** to return to the Enterprise View editor.

**Minimum/Maximum:** The minimum/maximum attributes function as user aids by providing defined limits for each column. For example, for the column *INCOME*, you could specify a minimum value of 20000 and a maximum value of 75000.

Both attributes can be edited by double-clicking in the cell and entering text manually.

**Description:** Table and column descriptions can be added or edited during the Add Table or Add Column tasks or by double-clicking within one of the *Description* column cells.

The Enterprise View Table/Column Description dialog box allows you to enter detailed table/column description text. The Enterprise View description is intended to provide high-level table and column information. The text defined here is also displayed (read-only) in the Application View Description dialog box.

## Managing keys

The Manage Keys dialog box provides options for viewing and defining keys for each table. A **key** is a set of columns that can be used to identify or access a particular row or rows. The key is identified in the description of a table, index, or referential constraint. The same column can be part of more than one key.

**Note:** A key can be based on more than one column. Rows in two tables match if the values of all columns in the table keys match each other, in the same order.

A unique key is a key that is constrained so that no two of its values are equal. The columns of a unique key cannot contain NULL values. For example, an employee number column can be defined as a unique key, because each value in the column identifies only one employee. No two employees can have the same employee number.

### Assigning keys to a table

1. Click **Manage Keys**. The Manage Keys dialog box opens.
2. Click **Add** to define a new key. The Key Properties dialog box displays, allowing you to define or edit an existing key.
3. Enter an appropriate key name in the **Name** field and click **Unique** if the key is unique. This name appears under the list of keys and is the name a consuming application will use to determine the key.
4. Select the appropriate columns from the Available Columns table and click the right-arrow button. The selected column names are displayed in the Key Columns list. Derived columns can be used as keys.  
You can remove existing columns from the Key Columns list by selecting the appropriate columns and clicking the left-arrow button.
5. Click the **Up** and **Down** buttons to rearrange the key column order.
6. Click **OK** to save the key definition. The Key Properties dialog box closes, and the key name is displayed in the Manage Keys dialog.
7. Click **Close**. The selected table displays the number of currently defined keys in the corresponding column.

### Manage links

The Manage Links dialog box provides options for viewing and defining links between tables and data sets.

Links are similar in function to SQL joins in that they allow you to use data from more than one data source table or data set. Tables and data sets can be related to each other with keys. A key is a column that contains a value for each row. Unique keys are defined for columns that contain unique values for each row (customer identification numbers for example). Refer to the “Managing keys” on page 209 section for more information on keys.

### Defining a link between two tables or data sets

1. Click **Manage Links**. The Manage Links dialog box opens.
2. Click **Add** to define a new link. The Link Properties dialog box displays, allowing you to define a new link.
3. Enter an appropriate link name in the **Name** field. This name appears under the **Name** column in the Manage Links dialog and is the name a consuming application will use to determine the link.
4. Select the appropriate table or data set and key under the **Left** section. The **Tables/Data Set** drop down menu provides a list of all defined tables or data sets. The **Key** drop-down menu provides a list of all keys currently defined for the selected table or data set. If the table is the target of the link, keys based on derived columns are omitted.
5. Select the appropriate table or data set and key under the **Right** section. The **Tables/Data Set** drop down menu provides a list of all defined tables or data sets. The **Key** drop-down menu provides a list of all keys currently defined for the selected table or data set. If the table is the target of the link, keys based on derived columns are omitted.

When matching on keys of the decimal type, the two keys must have the same scale. For example, the SQL Server *money* type is converted to a decimal with a scale of 4. A column of this type can only be linked to another column of the decimal type with a scale of 4.

The selected keys must have the same number of matching column types. For example, if the **Left** table's key contains five INTEGER column types, the **Right** table's key must also contain five INTEGER column types. Refer to the “Managing keys” on page 209 section for more information on keys. If



editing the keys used in a link results in incompatibilities between the keys, the link using those keys is removed from the system when you close the Link Properties dialog.

6. Select the appropriate **Allow Link** check boxes.
  - When both boxes are checked, the link maps freely between both tables/data sets.
  - When only **Right as Source, Left as Target** is checked, a single row in the **Right** table/data set maps to a row or rows in the **Left** table/data set.
  - When only **Left as Source, Right as Target** is checked, a single row in the **Left** table/data set maps to a row or rows in the **Right** table/data set.
  - When neither check box is checked, the link is deactivated.
7. Select the appropriate cardinality for both the **Left** and **Right** table/data set selections. The cardinality settings define the row relationships between the two selected tables/data sets. It is assumed you are starting from one row in the source table/data set, and defining how many rows in the target table/data set to map.

For example, assume:

- **Left** table/data set is defined as CUSTOMERS.
- **Right** table/data set is defined as CITY.
- **Left as Source, Right as Target** is defined as Exactly One (1..1).
- **Right as Source, Left as Target** is defined as Zero or More (0..\*).
- The **Allow Link** check box is selected for both tables/data sets.

Given this example, the following rules will be enforced:

- **Left as Source, Right as Target.** One customer maps to one city. Each customer must be registered in exactly one city.
- **Right as Source, Left as Target.** One city maps to any number of customers. Each city may have any number of customers, including none.

The values displayed in the **Cardinality** drop-down menu will vary depending on whether or not the selected key is unique.

When the selected key is unique, you are presented with the following **Cardinality** options:

- At Most One (0..1)
- Exactly One (1..1)

When the selected key is not unique, you are presented with the following **Cardinality** options:

- Zero or More (0..\*)
- At Most One (0..1)
- Exactly One (1..1)
- At Least One (1..\*)

8. Enter an appropriate alias for each link table/data set (optional). The alias defines a unique name for each link table/data set. Without the alias, the link name is used for each linked table/data set.
9. Click **OK** to save the link definition. The Link Properties dialog box closes, and the link name is displayed in the Manage Links dialog.
10. Click **Close** to return to the editor.

## Saving the Enterprise View

Saving the Enterprise View always creates a new version of the Enterprise View in the enterprise repository. It is not possible to overwrite previous versions. Object versions are handled by the built-in enterprise repository functionality.

To save a new version of the Enterprise View:

1. Click the save icon from the IBM SPSS Collaboration and Deployment Services Deployment Manager toolbar

or

2. Choose:

**File > Save**

All saved changes to the Enterprise View are logged and include the ID of the user who saved the changes, the Enterprise View version ID, and a timestamp reflecting when the changes were saved.

---

## Application View

An Application View collects a subset of tables and columns from the Enterprise View that relate to a specific application. It captures additional information about how tables and columns are used in the context of an application (for example, call center or fraud detection). There can be more than one Application View associated with the Enterprise View.

The Application View serves two purposes. First, it provides a means to constrain the information displayed to a user in a tool or application based on the IBM SPSS Collaboration and Deployment Services Enterprise View. For example, a user may be interested only in call center information. By identifying which IBM SPSS Collaboration and Deployment Services Enterprise View objects are used in the call center application, user interface logic can limit results to only those relevant objects. Second, the Application View provides a means for a system administrator or data expert to view the repository from the perspective of an application. This is especially useful in assessing the effect changes might have on a specific application.

## Creating an Application View

1. Right-click the *Content Repository* directory and choose:

**New > Application View...**

or

2. Choose:

**File > New > Application View...**

The New Application View dialog box opens, providing options for defining an Application View name and description. After entering an appropriate Application View name and description, click **Finish** to launch the "Application View editor."

## Application View editor

The Application View editor allows data experts to create, edit, review, and validate an Application View. An Application View is created from tables and columns defined in the Enterprise View definition. Consumers of an Application View, for example an IBM SPSS Modeler stream, can be linked to the Application View, facilitating impact analysis.

## Adding tables and columns to the Application View

Each Application View references a specific Enterprise View version. The available tables and columns are limited to what the Enterprise View version contains.

To specify which tables and columns are included in the Application View:

1. Select an appropriate table from the Application View editor's table section. Once a table is selected, its corresponding columns display in the column section.
2. Select columns by clicking the appropriate check boxes in the columns section.
3. Repeat the previous steps for each required table.

## Removing tables and columns from the Application View

To remove tables and columns from the Application View:

1. Select an appropriate table from the Application View editor table section. Once a table is selected, its corresponding columns display in the column section.
2. Remove columns by clearing the appropriate check boxes in the columns section. If all columns for a particular table are cleared, the table is automatically cleared in the table section.
3. Repeat the previous steps for each required table.

## Column count

The **Column Count** column indicates how many columns are currently selected for each specified table. For example, when the CUSTOMER table lists a column count of 4/20, that means that 4 out of 20 possible columns are currently selected. The column count quickly lets you verify the column coverage for each table.

## Sorting columns

All columns can be sorted in ascending or descending order. To sort a column, click an appropriate column heading from the column section. The arrow displayed in the heading indicates whether the column is sorted in ascending or descending order.

*Note:* Sorting does not affect how columns are displayed for consuming applications. The index order dictates how columns are displayed for consuming applications. See the topic “Index” for more information. .

## Changing Application View column attributes

Column attributes can be changed by simply double-clicking the column cell. You can also select multiple rows, right-click, and apply the same attribute changes to all selected rows (for example, you can apply the same **Environment** to all selected rows).

**Note:** Changing row attributes via multi-row selection erases the undo/redo history.

Not all Application View columns are user definable. For example, the **Column Name** and **Type** columns are defined in the Enterprise View editor and cannot be modified in the Application View editor.

- **Index:** The **Index** column represents the order in which Enterprise View columns will be presented to consuming applications. See the topic “Index” for more information.
- **Column Name:** Displays the current column name. This attribute cannot be modified in the Application View editor. See the topic “Column name” on page 214 for more information.
- **Type:** Specifies the type of data stored in a column. Each column can store data consisting of only a single data type. Refer to “Type” on page 214 for information regarding supported column types. This attribute cannot be modified in the Application View editor.
- **Environment:** Displays the environments currently associated with the column. See the topic “Environment” on page 215 for more information.
- **Direction:** Indicates whether the column is a predictor, a target, or both. See the topic “Direction” on page 215 for more information.
- **Description:** Provides a user-defined column description. See the topic “Description” on page 215 for more information.

**Index:** The column index order dictates how columns are displayed for consuming applications. In the following table, the columns will be displayed to consuming applications in the order *Income, State, Ownership, Mortgage, and Second Mortgage*.

Index	Column name
4	Mortgage
2	State
5	Second Mortgage
3	Ownership

Index	Column name
1	Income

By default, the index order is determined by the order in which the columns were defined in the Enterprise View. The index order can be edited manually at the Enterprise View or the Application View level.

*Note:* When the index order is changed at the Application View level, it does not affect the original Enterprise View indexing order.

**Column name:** The column name attribute displays the current column name inherited from the Enterprise View.

**Type:** The column type attribute is inherited from the Enterprise View and specifies the type of data stored in a column. Each column can store data consisting of only a single data type. See the topic “IBM SPSS Collaboration and Deployment Services Enterprise View data types” on page 254 for more information.

*Table 23. Data types*

Type	Description
Boolean	Used to store binary digits (0 or 1). These can be used to denote Boolean values, such as Yes/No, True/False, or On/Off.
Date	Values must have the form CCYY-MM-DD, where CCYY denotes the year, MM denotes the number of the month (January = 1, February = 2, and so on), and DD denotes the day (starting with 1). Optionally, the year can be preceded by + or -. Leading zeros for the year, month, and day are optional. Spaces are not allowed before or after the - separator.
Daytime	Values must have the form HH:MM:SS.xxx, where HH denotes the hours (ranging from 0 to 23), MM denotes the minutes (0 to 59), SS denotes the seconds (0 to 59), and xxx, the milliseconds (0 to 999). First, note that only the 24-hour representation is allowed. Furthermore, the seconds and milliseconds are optional. Leading zeros for the hours, minutes, and seconds are optional. Spaces are not allowed before or after the : separator.
Decimal	Used to store exact numbers. The scale and maximum precision can be configured. The precision denotes the maximum number of digits each value may consist of. The scale indicates how many of those digits may occur after the decimal point (that is, the fractional part). The value of this parameter must be greater than or equal to zero and less than or equal to the precision.
Double	Represents signed, approximate numerical values. Double permits much larger numbers than Float.
Float	Represents signed, approximate numerical values. Double permits much larger numbers than Float.
Integer	Used for whole numbers (for example, 5 or 110).
Long	Used for character data of variable length up to two gigabytes.
String	Used to store a sequence of ISO-8859-Latin-1 characters. The maximum length of such a sequence is configurable. Furthermore, the string may be empty and cannot contain the newline (ASCII 10) or null (ASCII 0) characters.
Timestamp	Allows values of the form CCYY-MM-DDTHH:MM:SS.xxx in accordance with the ISO 8601 standard. It is easy to see that it is nothing more than the concatenation of the Date and the Time types, using a delimiter of T to indicate the beginning of the Time portion. Consequently, the rules for the Date and the Time types are also applicable to the Timestamp type.

**Environment:** The environment settings provide a means of identifying which particular columns should be associated with defined business segments. The business segments include:

- **Analytical.** Columns associated with an Analytic environment typically define the transaction data needed to run analytic and optimization tasks.
- **Operational.** Columns associated with an Operational environment are typically used to deploy the resultant analytics to operational channels.
- **Reporting.** Columns associated with an Reporting environment are typically used to provide the resultant analytics to reporting channels.

The column environment determines the availability of the column in data provider definitions. The data provider definition is associated with a particular environment for the Application View. Only columns belonging to that environment will be available to the data provider definition. For example, an *Analytic* data provider definition has access only to columns in the associated Application View that are defined as *Analytical*.

Columns can be defined as belonging to multiple environments. For example, the column *CUST\_ID* can be assigned to both the *Analytical* and *Operational* roles. In this case, the column would be available in both *Analytic* and *Operational* data provider definitions.

*Note:* By default, the Application View editor lists all new columns as belonging to all three environments.

**Direction:** The direction setting specifies whether columns should be treated as predictors or targets. Predictors are used as an input to machine learning (predictor fields); targets are used as output for machine learning (predicted fields). The direction setting does not affect how columns are used or presented. The setting serves as a aid, allowing users to ascertain each column's intended role. IBM SPSS Modeler is one example of an application that utilizes the direction setting.

**Description:** Table and column descriptions can be added or edited by double-clicking inside one of the description column cells.

The Application View Description dialog displays any existing Enterprise View description information (read-only) and allows you to enter more detailed table/column description information that is specific to the Application View.

## Saving the Application View

You can save a new Application View to the repository or modify and save an existing Application View. Saving an existing Application View always creates a new version of the Application View in the repository. It is not possible to overwrite previous versions. Object versions are handled by the built-in repository functionality.

To save a new version of the Application View:

1. Click the save icon on the IBM SPSS Collaboration and Deployment Services Deployment Manager toolbar  
or
2. Select

**File > Save or Save As...**

All saved changes to the Application View are logged and include the ID of the user who saved the changes, the Application View version ID, and a timestamp reflecting when the changes were saved.

## Changing the referenced Enterprise View labeled version

By default, a new Application View uses the Enterprise View version labeled **LATEST** for defining table and column information. An existing Application View will first attempt to match its label with an equivalent Enterprise View label. If no label match is found, the Application View editor defaults to the Enterprise View version with the **LATEST** label.

*Note:* It is possible to associate another label with the **LATEST** version. For example, you could apply the label **PRODUCTION** to the latest Enterprise View version. In this case the most recent version contains both the **LATEST** and **PRODUCTION** labels.

In some instances you may not want to use the most current Enterprise View version.

1. Click **Change** to select a different Enterprise View labeled version on which the Application View is based.

*Note:* Changing to a different labeled version is only effective for the current editing session. If you want to permanently change the reference Enterprise View version you must ensure the Application View version label matches the Enterprise View version label. See the topic “Working with labels” on page 29 for more information.

See the topic “Working with labels” on page 29 for more information.

---

## Merging view versions

An Enterprise View or an Application View may have several different versions. Every saved change to an Enterprise View or Application View results in a new version of that view. Furthermore, when importing a view, all imported versions yield new versions of the corresponding target view.

Occasionally, information from two separate versions of a view is needed, such as a table from version *A* of a view and a table from version *B*. In this case, the versions can be merged into a new, composite version for the view. Note, however, that the merge involves two versions of the same view. Different views cannot be merged. For example, Application View *View1* cannot be merged with Application View *View2*, but version *A* of *View1* can be merged with version *B*.

To merge two versions of a view, right-click the view in the Content Explorer and select **Merge Versions**. The Select Versions to Merge dialog box appears.

**Preferred version.** Select the first version of the view to be merged. The tables, keys, links, and column information from the preferred version are preserved in the merged version, but table descriptions are not.

**Secondary version.** Select the version of the view to merge with the preferred version. Tables, keys, links, and columns from the secondary version are added to the preferred version. Merged tables contain everything from the preferred version with additional columns or metadata for existing columns from the secondary version. If a definition in the secondary version conflicts with a definition in the preferred version, the latter definition is used in the merged version by default. For example, if the preferred version defines the *Age* column of the *My\_Table* table as a *String* and the secondary version defines *Age* as an *Integer*, the merged version defines *Age* as a *String*.

**Merge conflicts individually.** When selected, the resolution for each conflict occurring when merging the selected versions can be manually specified individually instead of allowing the preferred version to automatically take precedence. See the topic “Resolve conflicts” on page 217 for more information.

After selecting the two versions, click **OK** to perform the merge.

Usually, the merged version contains definitions from either one of the two versions being merged, depending on whether conflict resolution is automatic or manual. However, in some cases, the result is not a strict selection from one of the two versions.

- When merging categorical values of an Enterprise View column, the set of categories for the merged column is the union of the categories from the two selected versions, ignoring case. To specify one of the category sets instead of the union, merge conflicts individually.
- If the two selected versions each have a key with the same name on the same table but with different defined properties, the key from the secondary version is renamed.
- Links are updated to use any new key names. However if a link cannot be included due to a changed column type in one of the keys, the link is not included in the merged version.

## Resolve conflicts

The Resolve Conflicts dialog box identifies conflicts arising when merging two versions of an Enterprise View or an Application View and defines a resolution policy for each table containing conflicts.

The table list identifies the tables in the selected views that exhibit conflicts.

- **Table Name.** Name of the table resulting in conflicts.
- **Resolution.** Specifies the conflict resolution policy for the table. The policy can use the values from either selected version of the view or define custom settings. To define the policy, select the cell and identify the approach from the list.

For a selected table, the column list displays the properties for columns having conflicts, with the properties available depending on whether the merge involves an Enterprise View or an Application View. For resolution policies based on a specific view version, the property values cannot be modified as one version is simply overriding the other. However, for a custom resolution policy, the values for conflicting properties may be selected individually. The conflicting property value are highlighted to assist in identifying which values must be specified.

**Preferred Values.** For a custom resolution policy, sets all column properties to the values defined for the preferred version of the view.

**Secondary Values.** For a custom resolution policy, sets all column properties to the values defined for the secondary version of the view.

To specify a value for an individual column property in a custom resolution policy, select the cell containing the value and identify the new value from the list.

After specifying the table resolution policies and any custom resolution values, click **Merge** to generate the new merged version of the view.

---

## Data Provider Definition

The Data Provider Definition is used exclusively for retrieving data from relational data sources in non-real time applications. The Data Provider Definition manages the data at each stage by mapping the logical column definitions in the Application View to physical table columns in the customer database. This may involve a direct mapping of columns or SQL queries to select or derive columns and records as needed. For example, the historic data used to build the model may come from the warehouse, and the data being scored may come from the operational system. The Data Provider Definition also specifies the data source and credentials used to access data.

A Data Provider Definition links source data directly to an Application View and indirectly to Enterprise View tables and columns. A single Application View table may have more than one Data Provider Definition, each of which may point to multiple data sources.

## Creating a Data Provider Definition

To create a Data Provider Definition:

- Right-click the **Content Repository** directory and select:  
**New > Data Provider Definition...**  
or
- Choose:  
**File > New > Data Provider Definition...**

The **New Data Provider Definition** wizard appears, providing options for defining a Data Provider Definition.

After entering an appropriate name and description for the Data Provider Definition, click **Next** to proceed to the “Select Application View” dialog box.

**Note:** The **Name** and **Description** fields are both limited to 256 characters.

### Select Application View

The **Select Application View** dialog provides options for associating the Data Provider Definition with an appropriate Application View. The dialog only lists views that currently exist in the repository.

1. Select an Application View and select the appropriate environment (Analytic, Operational, or Reporting). See the topic “Environment” on page 215 for more information.
2. Click **Next** to display the “Choose configuration objects” dialog.

### Choose configuration objects

The **Choose Configuration Objects** dialog provides options for selecting defined JDBC/ODBC data source definitions and server credentials. The dialog also provides options for creating new data source definitions and server credentials.

Note that both the JDBC and ODBC data sources must point to the same database (and they will have the same permissions because they share a credentials object). The reason for having both data source types is because the Data Provider Definition can consume both Java-based and C-based applications. For example, BIRT Report Designer for IBM SPSS is a Java-based application that can use the IBM SPSS Collaboration and Deployment Services Enterprise View to get information for its reports. Being Java-based means that it needs to use a JDBC driver to connect to the database. IBM SPSS Modeler, in contrast, can also use the IBM SPSS Collaboration and Deployment Services Enterprise View to retrieve data for its streams. But IBM SPSS Modeler is a C-based application, so it requires an ODBC connection to the database.

Refer to “Data source definitions” on page 204 and “Credential definitions” on page 204 for more information.

1. After selecting the appropriate data source definitions and credentials, click **Finish**. The new definition displays in the Data Provider Definition editor. See the topic “Data Provider Definition editor” for more information.

### Data Provider Definition editor

The Data Provider Definition editor allows data experts to create, edit, review, and validate a Data Provider Definition. A Data Provider Definition maps the column definitions in the Application View to physical tables in the customer database.



## Changing the referenced Application View labeled version

By default, a new Data Provider Definition uses the Application View version labeled **LATEST**. An existing Data Provider Definition will first attempt to match its label with an equivalent Application View label. If no label match is found, the Data Provider Definition editor defaults to the Application View version with the **LATEST** label.

*Note:* It is possible to associate another label with the **LATEST** version. For example, you could apply the label **PRODUCTION** to the latest Application View version. In this case the most recent version contains both the **LATEST** and **PRODUCTION** labels.

In some instances you may not want to use the most current Application View labeled version.

1. Click **Change** to select a different, labeled Application View version.

See the topic “Working with labels” on page 29 for more information.

## Mapping column definitions

The following steps explain the process of mapping the logical Application View tables and columns to physical data source tables and columns.

1. Click a table row that does not yet have mappings defined (the **Mappings** column is empty) to launch the Data Provider Table Specification dialog.

*Note:* You can also click inside a defined **Mapping** column cell. An ellipses button appears in the cell. Clicking the button launches the Data Provider Specification dialog.

The Data Provider Specification dialog provides options for selecting an existing data source table/view or manually requesting column and table information via an SQL statement.

2. When you select **Existing Table**, the drop-down field becomes active, allowing you to select an existing data source table or view:

After the appropriate table/view is selected, you can verify column information by clicking **Sample**. The Sample Data dialog displays all columns from the selected table/view and a fifty-row sampling of actual data.

3. Click **OK** to return to the Data Provider Definition editor. All columns from the selected table/view now display in the **Physical Columns** section.

*Note:* Only columns from a single table are returned when the **Existing Table** option is selected.

The Data Provider Definition editor will automatically map physical columns whose **Column Name** and **Type** match their logical column counterparts.

4. When you choose the **Select Statement** option, the edit box becomes active, allowing you to enter an SQL select statement. The select statement allows you to specify columns from more than one table/view or define various filtering attributes via a WHERE clause (WHERE GENDER = MALE, for example). The WHERE clause information is not used in the Enterprise View, but is applied at run time.

*Note:* The select statement must adhere to the syntax rules of the physical data source.

5. After entering an appropriate select statement, click **OK** to return to the Data Provider Definition editor. Only those columns defined in the select statement display, in the **Physical Columns** section, for the specified table(s).

The Data Provider Definition editor will automatically map physical columns whose **Column Name** and **Type** match their logical column counterparts.

## Manually mapping columns

1. Select a column name from the **Physical Data Source Columns** table.
2. Select a corresponding column name from the **Application View Columns** table.
3. Click **Map**. The Data Provider Definition editor will map columns if their **Type** is the same.

**Note:** Alternatively you can map columns by dragging and dropping the **Physical Data Source Columns** into the **Application View Columns** table. Again, the **Type** must match between columns.

### Unmapping columns

1. Select the appropriate columns from the **Application View Columns** table and click **Unmap**.

**Note:** Click **Unmap All** to remove the mappings for all columns.

### Validating the Data Provider Definition

A Data Provider Definition is valid only when all tables and columns are properly mapped. The Data Provider Definition editor provides the mapping status for each Application View table:

- A green check mark indicates that all logical table columns are correctly mapped.
  - A dotted box indicates that not all logical table columns are mapped.
  - A broken box indicates that the logical table contains columns that were once mapped, but are now broken. Broken mapping errors typically occur when the SQL or data source, upon which the Application View's tables and columns are derived, is changed.
  - A red exclamation with a blue arrow indicates that the logical column or table no longer exists in the Application View.
1. Click **Validate** to check the mappings for all Application View tables. When all tables/columns are properly mapped, the Validation Status dialog indicates that the Data Provider Definition is valid.

When there is a problem with the column mappings, the Validation Status dialog will display one to three problems:

### Unmapped objects

Unmapped objects errors occurs when columns that exist in the Application View are not mapped in the Data Provider Definition. All Data Provider Definition column references must be mapped in order for the definition to be considered valid.

To resolve unmapped object errors:

- Ensure all Application View columns are properly mapped to the appropriate physical data source columns. See the topic "Mapping column definitions" on page 219 for more information.

or

- Remove the unmapped column references from the associated Application View.

### Broken mappings

Broken mapping errors typically occur when the SQL or data source, upon which the Application View's tables and columns are derived, is changed. The Data Provider Definition mappings were correct at some point, but after changing the data source, the mappings are no longer valid.

To resolve broken mapping errors:

- Change the referenced data source back to the data source upon which the original mappings were established or ensure the new data source contains the same tables and columns as the previous data source.

### Objects do not exist in the Application View

Missing Application View object errors occur when columns are removed from the Application View after they have been mapped in the Data Provider Definition.

To resolve missing object errors:

- Restore the missing columns to the Application View. See the topic “Adding tables and columns to the Application View” on page 212 for more information.

or

- Click **Fix** to automatically remove the column references from the Data Provider Definition.

## Saving the Data Provider Definition

Saving the Data Provider Definition always creates a new version of the Data Provider Definition in the repository. It is not possible to overwrite previous versions. Object versions are handled by built-in repository functionality.

To save a new version of the Data Provider Definition:

1. Click the save icon on the IBM SPSS Collaboration and Deployment Services Deployment Manager toolbar
- or
2. Select

**File > Save or Save As...**

All saved changes to the Data Provider Definition are logged and include the ID of the user who saved the changes, the Data Provider Definition version ID, and a timestamp reflecting when the changes were saved.

---

## Data Provider Definition - Real Time

The Data Provider Definition - Real Time is used exclusively for real time interactions with applications. A Data Provider Definition - Real Time links data sets directly to an Application View and indirectly to Enterprise View tables and columns. A single Application View table may have more than one associated Data Provider Definition - Real Time, each of which may point to a different data source.

A data set represents a source of data. In its simplest form, a data set maps to a single data source table and abstracts the physical implementation of the data source. In addition to mapping to a data source table, a data set can map to other data sets. Using this mechanism, data from various data sources can be linked together.

The Data Provider Definition - Real Time retrieves data sets, one row at a time, from the following sources:

- JDBC data source defined in the Data Provider Definition - Real Time
- Application server defined data source
- User context data defined in the Data Provider Definition - Real Time
- User created data supplier (implemented externally)

*Note:* Additional data columns can be derived from existing columns.

Refer to “Data sources” on page 48 for more information.

## Creating a Data Provider Definition - Real Time

To create a Data Provider Definition - Real Time:

1. Right-click the **Content Repository** folder and select:  
**New > Data Provider Definition - Real Time...**
- or
2. Choose:

### File > New > Data Provider Definition - Real Time...

The **New Data Provider Definition - Real Time** wizard appears, providing options for defining a Data Provider Definition - Real Time.

3. After entering an appropriate Data Provider Definition - Real Time name and description, click **Next** to proceed to the "Select Application View" step.

**Note:** The **Name** and **Description** fields are both limited to 256 characters.

## Select Application View

The **Select Application View** dialog provides options for associating the Data Provider Definition - Real Time with an appropriate Application View. The dialog only lists views that currently exist in the repository.

1. Select an Application View.

*Note:* **Operational** is the only available environment. See the topic "Environment" on page 215 for more information.

2. Click **Finish**. The **Select Application View** dialog closes and the Data Provider Definition - Real Time editor displays.

## Data Provider Definition - Real Time editor

The Data Provider Definition - Real Time editor allows data experts to create, edit, and review a Data Provider Definition designed for real-time interactions. A Data Provider Definition - Real Time maps the column definitions in a Application View to data sets derived from the following sources:

- JDBC data source defined in the Data Provider Definition - Real Time
- Application server defined data source
- User context data defined in the Data Provider Definition - Real Time
- User created data supplier (implemented externally)

Refer to "Data sources" on page 48 for more information.

The Data Provider Definition - Real Time editor includes two tabs:

- **Mapping.** The Mapping tab displays the currently defined data set mappings and provides options for manually configuring mappings. See the topic "Mapping data set column definitions" for more information.
- **Data Sets.** The Data Sets tab displays the currently defined data sets and provides options for manually configuring data sets, managing links, and managing keys. See the topic "Working with data sets" on page 224 for more information.

## Mapping data set column definitions

The following steps explain the process of mapping the logical Application View tables and columns to real time data set tables and columns.

1. Click a table row that does not yet have mappings defined (the **Mapping** column is empty).
2. Select an existing mapping from the drop-down menu in the selected table's **Mapping** column.

*Note:* You will need to create a data set mapping if no mappings are currently defined. See the topic "Working with data sets" on page 224 for more information.

The **Data Set Columns** table automatically populates with the column names and types defined for the selected data set mapping. The Data Provider Definition - Real Time editor automatically maps columns when both the **Column Name** and **Type** match.

**Note:** A data set can only be mapped to at most one Application View table.

3. Data is cached on the server. You can specify a cache timeout value in seconds. Position your cursor in the Cache Timeout column and type a value. For example, if you specify 30 for the cache timeout value, the data cached for the given Application View table will be cleared every 30 seconds.

## Manually mapping columns

1. Select a column name from the **Data Set Columns** table.
2. Select a corresponding column name from the **Application View Columns** table.
3. Click **Map**. The Data Provider Definition - Real Time editor will map columns if their **Type** is the same.

*Note:* Alternatively you can map columns by dragging and dropping the **Data Set Columns** into the **Application View Columns** table. Again, the **Type** must match between columns.

4. Continue mapping until all columns in the **Application View Columns** table are properly mapped for each Application View table.
5. You can also specify whether the data in a given Application View column can be cached. Select **Yes** or **No** in the *Cachable* column.

## Unmapping columns

1. Select the appropriate columns from the **Application View Columns** table and click **Unmap**.

*Note:* Click **Unmap All** to remove the mappings for all columns.

## Previewing data

You can select a table and verify column information by clicking **Preview**. The Preview Data dialog displays all columns from the selected table/view, and allows you to preview data for the selected data set, and add, edit, import, export, and remove context values for data sets of type context data.

1. Click **OK** to return to the Data Provider Definition - Real Time editor.

**Previewing data:** The Preview Data dialog provides options for previewing data for the selected data set, and adding, editing, importing, exporting, and removing context values for data sets of type context data.

**Values for Key:** The drop-down list provides available defined keys. Keys based on derived columns are not available. Select an appropriate key. For some databases, the key type may impact the ability to preview the data.

- For SQL Server databases, the data cannot be previewed if the key is of type *text* or *ntext*.
- For DB2 databases, the data cannot be previewed if the key is of type *LONG VARCHAR*.
- For Netezza and Teradata databases, the data cannot be previewed if the key is of type *time* or *timestamp*.

**Add:** Opens the Value Editor dialog box, allowing you to specify a new value for the selected key.

**Edit:** Opens the Value Editor dialog box, allowing you to edit the selected key value.

**Import:** Opens the Import Values dialog box, allowing you to import key values from a *.csv* file.

**Export:** Opens the Export Values dialog box, allowing you to export key values to a *.csv* file.

**Remove:** Removes the selected key value.

**Remove All:** Removes all values for the selected key.

## Context Values

The Context Values section allows you to add, edit, import, export or remove values for context data type data sources.

**Add:** Opens the Value Editor dialog, allowing you to specify a new value for the context data columns.

**Edit:** Opens the Value Editor dialog, allowing you to edit context data columns values.

**Import:** Opens the Import Values dialog, allowing you to import context data columns values from a *.csv* file.

**Export:** Opens the Export Values dialog, allowing you to export context data columns values to a *.csv* file.

**Remove:** Removes the selected context data column values.

**Remove All:** Removes all context data column values.

**Get Data:** Retrieves context data column values from the data source.

## Working with data sets

The Data Sets tab provides options for defining, editing, and removing data sets.

A data set represents a source of data. In its simplest form, a data set maps to a single data source table and abstracts the physical implementation of the data source. In addition to mapping to a data source table, a data set can map to other data sets. Using this mechanism, data from various data sources can be linked together.

The Data Provider Definition - Real Time retrieves data sets, one row at a time, from the following sources:

- JDBC data source defined in the Data Provider Definition - Real Time
- Application server defined data source
- User context data defined in the Data Provider Definition - Real Time
- User created data supplier (implemented externally)

*Note:* Additional data columns can be derived from existing columns.

The system validates changes to the data source settings to determine if the modifications affect other definitions in the Data Provider Definition - Real Time. If the changes result in conflicts, such as missing necessary columns or incompatible data types, the modified data source cannot be applied to the data source. Other changes, such as the addition of columns, result in prompting for handling instructions. Refer to "Data sources" on page 48 for more information.

## Data sets

The **Data Sets** section provides options for defining data sets and working with data set links and keys.

**Add...** Launches the Add Data Set wizard, allowing you to define any number of data sets. See the topic "Defining data sets" on page 225 for more information.

**Manage Links.** Opens the Manage Links dialog box, allowing you to create new data set links or work with existing links. See the topic "Manage links" on page 210 for more information.

**Manage Keys.** Opens the Manage Keys dialog box, allowing you to create new keys or work with existing keys. See the topic "Managing keys" on page 209 for more information.

**Edit.** Updates the data set to include new columns added to the underlying source and to swap one underlying source for another provided the new source provides the same set of base columns. Removing base columns is not allowed once they have been added. See the topic "Editing data sets" on page 226 for more information.

**Remove.** Removes the currently selected data set.

## Columns

The **Columns** section provides options for previewing data set column data and working with derived columns.

**Preview.** Launches the Preview Data dialog box, allowing you to preview data for the selected data set, and add, edit, import, export, and remove context values for data sets of type context data.

**Add...** Launches the Create Derived Column expression editor, allowing you to manually create derived columns and join data sets. See the topic “Expression Editor” on page 227 for more information.

**Edit...** Launches the Edit Derived Column expression editor, allowing you to edit an existing derived column or data set join. See the topic “Expression Editor” on page 227 for more information.

**Remove.** Removes a selected derived column from the Columns list.

**Defining data sets:** The following steps explain the process of defining a new data set.

1. In the **Data Sets** section, click **Add** to launch the Add Data Set wizard.

The Add Data Set wizard's **Select data source** step allows you to define a JDBC, application server, or data service data source.

*JDBC data source:* The Data Provider Definition - Real Time editor uses the JDBC metadata to automatically import the primary and foreign keys that are defined on the SQL table or view. Using this information, links are automatically created between data sets.

*Note:* The interaction server must be running in order to change anything related to the JDBC data source (selecting a different table for example).

1. Select **JDBC** and click **Next** to display the Connection Settings step. For data sets that use a JDBC data source, the data set will map to a single SQL table (or view) within a SQL database.

*Connection settings for a JDBC data source:* The Connection Settings dialog box provides options for selecting a JDBC data source provider and data source credentials. The dialog also provides options for creating new data sources and credentials. Refer to “Selecting a data source definition type” on page 49 and “Adding new credentials” on page 47 for more information.

**JDBC.** Select an appropriate JDBC data source definition.

**Credentials.** Select an appropriate credential definition.

**Connect.** You must click **Connect** to verify the defined data source connection.

1. Click **Next** to select tables and columns.

*Select tables and columns for a JDBC data source:* The Select Tables and Columns dialog box provides options for selecting an existing data source table and its associated column definitions.

**Table Types...** Click to select what to return from the physical data source.

- **System Table**—Returns data source system tables (for example, sysdevices and sysfiles).
- **Table**—Returns regular data source tables (for example, CUSTOMERS and ORDERS).
- **View**—Returns table views. Views are virtual tables whose contents are defined by a query. As with a real table, a view consists of a set of fields and records. The records and fields come from tables referenced in the query defining the view and are produced dynamically when the view is referenced.

*Note:* Other options may be available, depending on the data source configuration.

1. Select the tables and columns and click **Finish**. You are returned to the Data Sets tab.

The newly defined data set, and its associated column definitions, are added to the **Data Set** table

*Application server data source:* For data sets that use an Application Server data source, the data set will map to a single SQL table or view within a SQL database. You do not specify a credential when connecting to the data source. The credentials are defined in the application server.

The data source is available in the application server where the interaction server is running. As a result, the Data Provider Definition - Real Time editor connects to the interaction server in order to retrieve data source metadata. The Data Provider Definition - Real Time editor uses the metadata to automatically import the primary and foreign keys that are defined on the SQL table or view. Using this information, links are automatically created between data sets.

1. Select **Application Server** and select the appropriate application server data source definition from the drop-down list.

Refer to “Data sources” on page 48 for more information.

2. Click **Next** to display the Select Tables to Add step.

*Data service data source:* For data sets that use a Data Services data source, the data set will map to a single table that is defined in the data service. The Data Provider Definition - Real Time editor retrieves data source metadata (columns, column types, and so on) using the information defined and stored as part of the repository object. No connection to an actual server is made.

*Note:* There can only be one Data Services data source of type **Context Data** per Data Provider Definition - Real Time.

1. Select **Data Service** and select the appropriate real-time data service data source definition from the drop-down list.

Refer to “Data sources” on page 48 for more information.

2. Click **Next** to display the “Select data sets to add for a data service data source” step.

*Select data sets to add for a data service data source:* The Add Data Set wizard's **Select data source to add** step allows you to select which data set to use from the previously selected real-time data service data source definition.

1. Select the appropriate data sets and click **Finish** to return to the Data Sets tab. The newly defined data set, and its associated column definitions, are added to the **Data Set** table.

### **Editing data sets:**

To edit a data set:

1. Select the data set from the **Data Sets** section.
2. Click the **Edit** button.
3. Specify the data source settings. See the topic “Defining data sets” on page 225 for more information. For JDBC data sources, specify the connection information.
4. Select the columns to include in the data set.

For JDBC and Application Server data sources, the **Table Types** button filters the list of available tables to selected types. Select a table from the list to view a list of the table columns. Each column will have one of the following states:

- **Missing.** The column appears in the data set being edited but the selected table does not contain a column having the same name.
- **Wrong type.** The column appears in both the data set and the selected table, but the types do not match.



- **Matched.** The column appears in both the data set and the selected table, and the type in each is identical. The column is selected for inclusion and cannot be cleared.
- **Additional.** The column appears only in the selected table but can be selected for inclusion in the data set.

If any columns have a state of *Missing* or *Wrong type*, the selected table cannot be used as the new data source.

After selecting columns, click the **Finish** button to commit the columns to the data set.

## Saving the Data Provider Definition - Real Time

Saving the Data Provider Definition - Real Time always creates in a new version of the Data Provider Definition in the repository. It is not possible to overwrite previous versions. Object versions are handled by the built-in repository functionality.

To save a new version of the Data Provider Definition - Real Time:

1. Click the save icon on the IBM SPSS Collaboration and Deployment Services Deployment Manager toolbar  
or
2. Select

**File > Save or Save As...**

All saved changes to the Data Provider Definition - Real Time are logged and include the ID of the user who saved the changes, the Data Provider Definition - Real Time version ID, and a timestamp reflecting when the changes were saved.

---

## Expression Editor

The expression editor allows the creation of derived attributes for a Data Provider Definition used with the IBM SPSS Collaboration and Deployment Services Enterprise View. The expression editor interface includes the following components:

- **Expression Panel.** Contains the expression canvas and controls for creating and editing expressions.
- **Schema Browser.** Allows you to browse through the schema and copy attributes to the expression.
- **Functions Panel.** Contains tabs with buttons for all available functions.

## Expression Panel

The expression panel includes the following:

**Name** specifies the name of the expression. If creating a new expression, enter a unique name of your choice.

**Type** indicates the available return types based on the application. If creating a new expression, select a return type from the menu. The selected return type must match the return type of the expression created.

**Expression Status** indicates whether the syntax of the expression is valid. If the syntax is valid, a status message indicating the return type is displayed. If the syntax is not valid, an error message is displayed.

**Show Boxes** indicates whether sub-expressions should be outlined in boxes for easier viewing.

**Sub-expression Name** specifies the name of a sub-expression used in the expression.

**Clear Selection** deletes the selected item.

**Clear All** clears the expression canvas.

**Preview** displays the result of the expression.

## Annotating Expressions

In the **Sub-expression Name** field, you can enter a name for a part of an expression. By giving one or more parts of an expression a useful name, complicated expressions can be made more easily readable because you can collapse and expand those named parts of expressions. The **Show Boxes** check box displays boxes in the expression canvas, illustrating which parts of the expression belong together.

## Previewing an Expression

You can preview the result of an expression by clicking the **Preview** button. You can also preview the result of an expression or subexpression by right-clicking it on the expression canvas and choosing **Preview of selection**. Always inspect the results of an expression before creating a new attribute.

## Cutting, Copying, and Pasting an Expression

Expressions can be cut, copied, and pasted into a text editor or another expression editor session running on the same machine. Simply select the expression, right-click, and choose **Cut** or **Copy** from the menu. Then paste the expression into the application.

Note that when you paste an expression from expression editor into a text editor, that expression is displayed in XML format. To add the expression back into expression editor, you must copy the entire XML block and paste it onto the expression canvas. The expression editor converts the XML back into an easy-to-read string representation of the expression.

## Functions Panel

Each tab on the functions panel contains a group of functions related to one another.

- The functions on the **Aggregate** tab can combine all the records of one table with the linked records from another table and return just one result, which can subsequently be used as the value for a new attribute in the first table. *Note:* Making use of aggregate functions in an expression can have a highly negative impact on the performance of its evaluation.
- The **Arithmetic** tab contains basic arithmetic functions you can use with numbers.
- The functions of the **Conversion** tab convert values from one data type to another.
- The **Date & Time** tab contains date and time functions.
- The **Miscellaneous** tab contains functions not offered on other tabs.
- The **Selection** tab contains functions to compare the values of expressions (like or, between, and isnil).
- The **String** tab contains string functions.

## About Functions

IBM Corp. functions perform many different kinds of operations. Functions are used to manipulate and prepare data. Thus, the raw data contained in tables (data sets) can be examined and extrapolated, providing more sophisticated information about customers and their attributes.

Before providing an overview of the different functions, this chapter explains:

- The different types of functions
- Data types associated with the functions
- Handling nil values
- Inspecting the results of an expression

**Types of Functions:** The expression editor includes two main types of functions: basic and aggregate.

**Basic functions** are functions used to derive results from the values of one or more attributes *from the same table*. For example, the string functions for string operations, the mathematical functions for mathematical operations, and the conversion functions to convert values into other types of values.

**Aggregate functions** are placed in the Aggregate tab. They can *combine* all records of one table with the linked records from another table and return one result. This result can be used as the value for a new attribute in the first table. For example, you can combine all records of the *PORDER* table that are linked to a specific record of the *Customer* table.

**Data types:** The input and return types for each function are listed throughout this guide. You can also hover over a function in the expression canvas to display a message indicating the valid data types. If you create an expression where one or more input values is not of the appropriate data type, the **Expression Status** field displays an error message. Similarly, if the return type of the expression does not match the return type listed in the **Type** field, an error is displayed when you attempt to add the new expression.

Where **Any** is listed as the data type, the input or return can be any single value of any data type. Where **Any:1** is listed, the input or return can be one *or more* values of any data type.

The Conversion tab contains functions to convert values from one data type to another. For more information, see “Conversion Functions” on page 237.

**Handling Nil Values:** The internal database of the server has a representation for unknown values. These are called **nil values** or **nils**.

When creating an expression, consider how nil values will be handled. The `ifnil` and `isnil` functions on the Selection tab are useful for testing for and excluding nil values. You can also use the `ifnil` function on the Miscellaneous tab to replace a nil value with another value.

Direct comparison (using `=`) of any value with nil will always give nil, even when the value being compared actually is nil. Therefore, selections such as `[Name] = str(nil)` will return the empty set. Always use `isnil([Name])` instead.

This section explains how the server deals with nil values in various situations.

## Basic Functions

Most basic functions (row-wise operations) will result in nil if any of the values they operate on are nil. For example:

- The `isnil` function returns true if the argument is nil and false if not.
- The `ifnil` function returns the second argument if the first argument is nil, otherwise it returns the first argument.

The following examples illustrate the use of the functions:

- `isnil(nil) → true`
- `isnil(3.0) → false`
- `[Total] = nil → nil` (direct comparison with nil always gives nil)
- `isnil([Total]) → false` (if `Total <> nil`) or `true` (if `Total = nil`)
- `ifnil(nil,9) → 9`
- `ifnil(8,9) → 8`
- `ifnil([Total],1000) → Total` (if `Total <> nil`) or `1000` (if `Total = nil`)

Certain functions on the Selection tab (`or`, `and`, `implies`) do not return nil under particular conditions. For example:

false and nil → false

In this case, the result will be false because one of the values is false. Similarly:

nil or false or true → true

Here, one of the values is true, so the result must be true.

### When Function

When a nil is matched against the branches of a when functor, the otherwise branch will not be taken. This is because nil is an unknown value so it is not known which branch is the correct one to take. Instead, the result is nil.

The following example illustrates the *wrong* way to deal with nils; it will return nil when [GENDER] is nil, instead of unknown as was intended:

```
"female" When: ([GENDER] = "F")
```

```
"male" When: ([GENDER] = "M")
```

```
"unknown" Otherwise.
```

This is the correct way to formulate this expression:

```
"female" When: ([GENDER] = "F")
```

```
ifnil("male" When: ([GENDER] = "M"), "unknown")
```

```
str( nil) Otherwise.
```

For more information about the when function, see “**when**” on page 244.

### Aggregate Functions

When using aggregate functions (column wise operation returning a unit value), if one of the values in the column is nil, the result will be nil.

*Exceptions:* The size and exist functions treat nil as a normal value.

The following examples illustrate the use of the functions:

```
min([Total]) → nil (if one of the values = nil)
```

```
min([Total]) → min-value (if no value = nil)
```

```
min([Total]) → in the Item table, nil will be returned
```

```
min(ifnil([Total],1000)) → in the Item table, min-value 52,10 will be returned because nil is replaced with 1000
```

```
size([Total]>=100) → in the Item table, 7 will be returned, because nil is treated as a normal value (< 100)
```

## Grouping

For grouping purposes, nils are considered equal and are therefore placed in their own group. This gives a better picture of the data than the alternate approach of considering every nil to be unique. For example:

`even_height(age,10)` → All nil values of age form a separate group.

Nils are also placed in their own group in the following situations:

- Finding the cardinality of a set using `card`
- Sampling an expression with unique tuples checked

## Links

When records in a table are matched to records in another table according to the values of the attributes in their keys, nils are considered distinct and do not match with other nil values. Similarly, if any of the attributes in a multi-attribute key is nil, no match is made.

## Data Upload

All values that do not belong to the domain return nil.

The following examples illustrate the use of the functions:

- `flt("3.215")` → 3.215
- `flt("a23")` → nil
- `flt("23.4a")` → nil
- `flt("")` → nil

This also applies when null values are uploaded from a remote database—such values are uploaded as nil values.

## Data Download

When data is downloaded to a database, nil values are returned null in the database. This means that attributes that are part of the key in the remote database cannot be nil (because keys in databases cannot be null).

**Results of an Expression in the Expression Editor:** Always inspect the results of an expression before creating a new attribute.

- Use the **Preview** button to inspect the entire expression, including all sub-expressions.
- If you want to inspect the contents of part of the expression, or of an attribute used in the expression, simply right-click the sub-expression or attribute and choose **Preview of selection**.

## Aggregate Functions

The Aggregate tab can be used to deal with 1–*n* links. For example, the link between the tables *CUSTOMER* and *ORDER* is a 1–*n* link as 1 customer can have placed 0–*n* orders. When deriving new attributes for the *CUSTOMER* table, set-valued expressions must be reduced to a single value. This is what the aggregate functions do; they take a set of values as input and produce a single value as output.

Thus, the functions of the Aggregate tab can combine all the records of the order table that are linked to a specific *CUSTOMER\_ID* and return just one result per customer. This result can be used as the value for a new attribute in the *Customer* table.

## Restriction:

Aggregate functions allow you to add a restriction in a link to restrict the result of an aggregate expression. The result type of the restriction must be boolean.

*Example:* You are only interested in counting customer orders with *Commissionplan Y*. Using a restriction (in this case *Commissionplan = Y*) means that, instead of considering all the records of the *Porder* table that are linked to a specific *CUSTOMER\_ID*, only the records for which the restriction is true will be considered.

**Note:** A link is a kind of restriction because, via a link, only the records linked to one specific record of another table will be considered (for example, the order records of a specific customer).

**avg:** The avg function returns the average of an expression.

### Syntax

Input Type	Return Type
avg(Integer)	Integer
avg(Long)	Long
avg(Float)	Float
avg(Double)	Double

The return type of the avg function is the same as the return type of the expression.

### Handling Nil Values

The avg function returns nil if the value of one of the attributes of the expression is nil.

### Example:

**count:** The count function returns the number of records in a linked table that are linked to the current record (without inspecting the values of one of the attributes). The function returns 0 if there is no linked record.

### Syntax

Input Type	Return Type
count(Any)	Integer

The return type is an integer (0, 1, 2, etc.).

### Handling Nil Values

Nil values do not apply to the count function—the function does not inspect attribute values.

### Example 1:

### Example 2:

**exist:** The exist function checks whether at least one of a set of booleans is true. The function returns false when there is no linked record.

## Syntax

### Input Type

exist(Boolean)

### Return Type

Boolean

The return type is a bit with the values false or true.

## Handling Nil Values

exist{true, nil} returns true

exist{false, nil} returns false

exist{nil, nil} returns false

**max:** The max function returns the maximum value of an expression. It returns nil if there is no linked record.

## Syntax

### Input Type

max(Any:1)

### Return Type

Any:1

The return type of the max function is the same as the return type of the expression.

## Handling Nil Values

The max function returns nil if the value of one of the attributes of the expression is nil.

### Example 1:

### Example 2:

**min:** The min function returns the minimum value of an expression. It returns nil if there is no linked record.

## Syntax

### Input Type

min(Any:1)

### Return Type

Any:1

The return type of the min function is the same as the return type of the expression.

## Handling Nil Values

The min function returns nil if the value of one of the attributes of the expression is nil.

**size:** The size function returns the number of true values of an expression (that is, the number of records for which the expression is known to be true). The function returns 0 if there is no linked record.

## Syntax

### Input Type

size(Boolean)

### Return Type

Integer

## Handling Nil Values

The size function considers nil to be a normal value and does not return nil if the value of one of the attributes of the expression is nil.

**sum:** The sum function returns the sum of an expression. It returns 0 if there is no linked record.

## Syntax

### Input Type

sum(Integer)

sum(Long)

sum(Float)

sum(Double)

### Return Type

Integer

Long

Float

Double

The return value is of the same type as those of the summarized attributes.

## Handling Nil Values

The sum function returns nil when the value of one of the attributes of the expression is nil.

## Arithmetic Functions

The Arithmetic tab contains basic arithmetic functions that can be used with all types of numbers (for example, integers, longs, floats, doubles, and decimals).

**Modulo:** %

Modulo divides one integer by another and returns the remainder.

## Syntax

### Input Type

Integer % Integer

### Return Type

Integer

## Example:

3 % 2 → 1

**Multiplication:** \*

The \* function is a regular arithmetic operator and can be used with all numeric types. The return value has the same type as the expression.

## Syntax

### Input Type

Integer \* Integer

### Return Type

Integer



<b>Input Type</b>	<b>Return Type</b>
Long * Long	Long
Float * Float	Float
Double * Double	Double
Decimal * Decimal	Decimal

**Example:**

12.13 \* 4 → 48.52

**Addition: +**

The + function is a regular arithmetic operator and can be used with all numeric types. The return value has the same type as the expression.

**Syntax**

<b>Input Type</b>	<b>Return Type</b>
Integer + Integer	Integer
Long + Long	Long
Float + Float	Float
Double + Double	Double
Decimal + Decimal	Decimal

**Example:**

12.13 + 4 → 16.13

**Subtraction: -**

The - function is a regular arithmetic operator and can be used with all numeric types. The return value has the same type as the expression.

**Syntax**

<b>Input Type</b>	<b>Return Type</b>
Integer - Integer	Integer
Long - Long	Long
Float - Float	Float
Double - Double	Double
Decimal - Decimal	Decimal

**Example:**

12.13 - 4 → 8.13

**Negation: -**

The - function is a regular arithmetic operator and can be used with all numeric types. The return value has the same type as the expression. -() is the negation operator.

### Syntax

Input Type	Return Type
-(Integer)	Integer
-(Long)	Long
-(Float)	Float
-(Double)	Double
-(Decimal)	Decimal

### Examples:

-(12.13)	→	-12.13
-(-15)	→	15

### Division: /

The / function is a regular arithmetic operator and can be used with all numeric types. The return value has the same type as the expression.

### Syntax

Input Type	Return Type
Integer / Integer	Integer
Long / Long	Long
Float / Float	Float
Double / Double	Double
Decimal / Decimal	Decimal

### Examples:

12.13 / 4	→	3.0325
3 / 4	→	0

**abs:** The abs function returns the absolute value of an expression (a number without its sign). This function can be used with all numeric types. The return value has the same type as the expression.

### Syntax

Input Type	Return Type
abs(Integer)	Integer
abs(Long)	Long
abs(Float)	Float
abs(Double)	Double
abs(Decimal)	Decimal

### Examples:

abs(12.13)	→	12.13
abs(-15)	→	15

## Conversion Functions

The functions of the Conversion tab convert a value of any type into the specified type.

### Handling Nil Values

When one of the expressions has a value of nil or cannot be converted, the functions of the Conversion tab return nil.

**bit:** The bit function converts an integer or string into a boolean value (logical true or false). If the value starts with one of {0, f, F, n, N}, it returns false; all other values return true.

### Syntax

Input Type	Return Type
bit(Integer)	Boolean
bit(String)	Boolean

### Examples:

bit(0)	→	false
bit("B")	→	true

**date (Conversion):** The date function converts a string containing a date into a date. The string must have the format *yyyy-mm-dd*, where *yyyy* denotes the year, *mm* the month, and *dd* the day.

### Syntax

Input Type	Return Type
date(String)	Date

### Example:

date("2008-10-31") → 2008-10-31

**dbl:** The dbl function converts an expression into a double. This function is not applicable to string expressions containing non-numeric characters.

*Note:* The decimal separator must be a decimal point. If not, the expression is not valid.

### Syntax

Input Type	Return Type
dbl(Integer)	Double
dbl(Long)	Double
dbl(Float)	Double
dbl(String)	Double

Input Type	Return Type
dbl(Decimal)	Double

**Examples:**

dbl(12.674567)	→	12.6746
dbl("90089")	→	90089

**decimal:** The decimal function converts a number of any type into a decimal. The decimal separator must be a decimal point. If not, the expression is not valid. The function is not applicable to string expressions containing non-numeric characters.

**Syntax**

Input Type	Return Type
decimal(Integer)	Decimal
decimal(Long)	Decimal
decimal(Float)	Decimal
decimal(String)	Decimal
decimal(Double)	Decimal

**Examples:**

decimal("90089.56")	→	90089.56
---------------------	---	----------

**flt:** The flt function converts an expression of any type into a float. The decimal separator must be a decimal point. If not, the expression is not valid. The function is not applicable to string expressions containing non-numeric characters.

**Syntax**

Input Type	Return Type
flt(Integer)	Float
flt(Long)	Float
flt(Double)	Float
flt(String)	Float
flt(Decimal)	Float

**Examples:**

flt(12)	→	12.0
flt("90089")	→	90089.0

**int:** The int function converts an expression of any type into an integer (whole number). The function is not applicable to string expressions containing non-numeric characters. The integer conversion of true returns 1. The integer conversion of false returns 0.

## Syntax

Input Type	Return Type
int(Boolean)	Integer
int(Long)	Integer
int(Float)	Integer
int(Double)	Integer
int(String)	Integer
int(Decimal)	Integer

## Examples:

```
int(12.67)           →      12
int(true)            →      1
```

**Int:** The `Int` function converts an expression of any type into a long (large whole number). The function is not applicable to string expressions containing non-numeric characters.

## Syntax

Input Type	Return Type
lng(Integer)	Long
lng(Float)	Long
lng(Double)	Long
lng(String)	Long
lng(Decimal)	Long

## Examples:

```
lng(12.67)           →      12
lng("325678")        →      325678
```

**str:** The `str` function converts an expression of any type into a string.

## Syntax

Input Type	Return Type
str(Any)	String
str(Decimal, String)	String

## Examples:

```
str(123)             →      "123"
str(true)            →      "true"
```

**timestamp:** The timestamp function converts a string containing a timestamp into a timestamp. The string must have the format *yyyy-mm-dd hh:mm:ss.fff*, where *yyyy* denotes the year, *mm* the month, *dd* the day, *hh* the hours, *mm* the minutes, *ss* the seconds, and *fff* the milliseconds.

### Syntax

Input Type	Return Type
timestamp(String)	Timestamp

### Example:

```
timestamp[("2004-04-18 09:23:45 987")] → 2004-04-18 09:23:45 987
```

## Date and Time Functions

The Date and Time tab contains the functionality to deal with dates and times. The tab relates to the date, daytime, and timestamp data types.

**adddays:** The `adddays` function adds a specified number of days to a specified date and returns a new date value. It takes two arguments: the original date, and the number of days to add to that date. When the second argument is negative, it is subtracted from the first argument.

### Syntax

Input Type	Return Type
adddays(Date, Integer)	Date

### Handling Nil Values

If one of the parameters (date or days) is equal to nil, then the result is equal to nil.

### Example:

```
adddays(date(1972, 7, 14), 11030) → 2002-9-25
```

**addmonths:** The `addmonths` function adds a specified number of months to a specified date and returns a new date value. It takes two arguments: the original date, and the number of months to add to that date. When the second argument is negative, it is subtracted from the first argument.

### Syntax

Input Type	Return Type
addmonths(Date, Integer)	Date

### Handling Nil Values

If one of the parameters (date or days) is equal to nil, then the result is equal to nil.

### Example:

```
addmonths(2001-04-18, 11) → 2002-03-11
```

**addyears:** The `addyears` function adds a specified number of years to a specified date and returns a new date value. It takes two arguments: the original date, and the number of years to add to that date. When the second argument is negative, it is subtracted from the first argument.

## Syntax

### Input Type

addyears(Date, Integer)

### Return Type

Date

## Handling Nil Values

If one of the parameters is equal to nil, then the result is equal to nil.

### Example:

addyears(2001-04-18, 2) → 2003-04-18

**date:** The date function returns a date from a timestamp, or three integers representing the year, month, and day. In the latter case, the first argument is year, the second argument is month, and the third argument is day.

## Syntax

### Input Type

date(Timestamp)

date(Integer, Integer, Integer)

### Return Type

Date

Date

## Handling Nil Values

- If the timestamp is equal to nil, then the function returns nil.
- If the year, month, or day is equal to nil or out of range (for example, day is equal to 32), then the function returns nil.

### Examples:

date(2001-04-18 09:23:15:965)

→

2001-04-18

date(2001, 04, 18)

→

2001-04-18

**day:** The day function returns the day of the month (value between 1 and 31) of a specified date.

## Syntax

### Input Type

day(Date)

### Return Type

Integer

## Handling Nil Values

If the date parameter is equal to nil, the function returns nil.

### Example:

date(2004-04-18) → 18

**dayofweek:** The dayofweek function returns the day of the week (value between 1 and 7) of a specified date. Sunday is equal to 1, Monday is equal to 2, and so on.

## Syntax

### Input Type

dayofweek(Date)

### Return Type

Integer

## Handling Nil Values

If the parameter date is equal to nil, the function returns nil.

### Example:

dayofweek(2001-04-18) → 4(Wednesday)

**dayofyear:** The dayofyear function returns the day of the year of a specified date. The first of January is equal to 1, the second of January is equal to 2, etc.

## Syntax

### Input Type

dayofyear(Date)

### Return Type

Integer

## Handling Nil Values

If the date parameter is equal to nil, the function returns nil.

### Example:

dayofyear(2001-04-18) → 108

**diff:** The diff function returns the number of days between two specified dates. If the first date is before the second date, the result is negative.

## Syntax

### Input Type

diff(Date, Date)

### Return Type

Integer

## Handling Nil Values

If one of the date parameters is equal to nil, the function returns nil.

### Examples:

diff(2001-04-18, 2001-04-16) → 2

diff(2001-04-18, 2001-04-20) → -2

**month:** The month function returns the month (value between 1 and 12) of a specified date.

## Syntax

### Input Type

month(Date)

### Return Type

Integer



## Handling Nil Values

If the date parameter is equal to nil, this function returns nil.

### Example:

```
month(2001-04-18) → 4
```

**weekofyear:** The weekofyear function returns the week of the year from the date. A new week starts with Sunday.

### Syntax

Input Type	Return Type
weekofyear(Date)	Integer

## Handling Nil Values

If the date parameter is equal to nil, the function returns nil.

### Example:

```
weekofyear(2001-04-18) → 16
```

**year:** The year function returns the year of a specified date.

### Syntax

Input Type	Return Type
year(Date)	Integer

## Handling Nil Values

If the date parameter is equal to nil, the function returns nil.

### Example:

```
year(2001-04-18) → 2001
```

## Miscellaneous Functions

The Miscellaneous tab contains functions that do not belong on other tabs.

## Handling Nil Values

Where one of the expressions has value nil, the functions of the Miscellaneous tab return nil.

*Exception:* The ifnil function replaces nil with a new value.

**ifnil:** The ifnil function replaces each nil value with a new value. It takes two arguments: the original value of the expression, and the new value to replace each nil value. If the value of the expression is not nil, then the original value of the expression is returned in the new expression.

*Note:* The first argument cannot be a boolean.

## Syntax

### Input Type

ifnil(Any:1, Any:1)

### Return Type

Any:1

## Handling Nil Values

The ifnil function replaces nil with a new value.

### Examples:

ifnil("hello", str(nil))	→	"hello"
ifnil(str(nil), "see you")	→	"see you"
ifnil("goodbye", "see you")	→	"goodbye"
ifnil(str(nil), str(nil))	→	nil

**when:** The when function checks a specified condition and returns a value as appropriate. It takes three arguments: the condition to check, the value to return when the condition is true, and the value to return when the condition is false.

In the expression editor, the function when is displayed as value-when-true when the condition is met, and value-when-false otherwise.

You can nest a when function in another when function. If you press Enter after the condition, a new when function is added.

## Syntax

Any:1 When: Boolean

Any:1 Otherwise.

## Handling Nil Values

The when function returns nil if the condition returns nil. For additional details about how the when function handles nil values, see “Handling Nil Values” on page 229.

### Example 1:

"hello"	When: true	→	"hello"
"goodbye"	Otherwise.		

### Example 2:

"hello"	When: false	→	"goodbye"
"goodbye"	Otherwise.		

### Example 3:

"hello"	When: false	→	"see you"
"see you"	When: true		
"farewell"	When: false		
"goodbye"	Otherwise.		

## Selection Functions

The Selection tab contains functions for comparing the values of expressions.

*Note:* All operators in the Selection operator group return a boolean value.

### Not Equal To: !=

The != function compares two expressions and returns true if they are not equal.

#### Syntax

Input Type	Return Type
Any:1 != Any:1	Boolean

### Handling Nil Values

The != function returns nil if one of the expressions is nil.

#### Examples:

1 != 2	→ true
1 != 1	→ false

### Less Than: <

The < function compares two expressions and returns true if the first expression is less than the second expression.

#### Syntax

Input Type	Return Type
Any:1 < Any:1	Boolean

### Handling Nil Values

The < function returns nil if one of the expressions is nil.

#### Examples:

1 < 2	→ true
-------	--------

1 < 1                                   → false

### Less Than or Equal To: <=

The <= function compares two expressions and returns true if the first expression is less than or equal to the second expression.

#### Syntax

Input Type	Return Type
Any:1 <= Any:1	Boolean

### Handling Nil Values

The <= function returns nil if one of the expressions is nil.

#### Examples:

1 <= 2	→ true
2 <= 2	→ true
3 <= 2	→ false

### Equal To: =

The = function compares two expressions and returns true if they are equal.

#### Syntax

Input Type	Return Type
Any:1 = Any:1	Boolean

### Handling Nil Values

The = function returns nil if one of the expressions is nil.

#### Examples:

1 = 1	→ true
1 = 2	→ false

### Greater Than: >

The > function compares two expressions and returns true if the first expression is greater than the second expression.

#### Syntax

Input Type	Return Type
Any:1 > Any:1	Boolean

## Handling Nil Values

The > function returns nil if one of the expressions is nil.

### Examples:

```
2 > 1           → true
1 > 1           → false
```

## Greater Than or Equal To: >=

The >= function compares two expressions and returns true if the first expression is greater than or equal to the second expression.

### Syntax

Input Type	Return Type
Any:1 >= Any:1	Boolean

## Handling Nil Values

The >= function returns nil if one of the expressions is nil.

### Examples:

```
1 >= 2          → false
2 >= 2          → true
2 >= 1          → true
```

**and:** The and function checks whether two boolean expressions are both true.

### Syntax

Input Type	Return Type
Boolean and Boolean	Boolean

## Handling Nil Values

If one of the boolean expressions returns nil, the result depends on the other boolean expression. If that expression returns true, the result is nil. If that expression returns false, the result is false.

### Examples:

```
true and true   → true
true and nil    → nil
nil and true    → nil
false and nil   → false
nil and false   → false
```

**between:** The `between` function compares three expressions and determines whether the value of the first expression is at least equal to the value of the second expression, and at most equal to the value of the third expression. So the second expression serves as lower boundary and the third as upper boundary for the first expression.

### Syntax

#### Input Type

`between(Any:1, Any:1, Any:1)`

#### Return Type

Boolean

### Handling Nil Values

The `between` function returns `nil` when the first expression is `nil`. A `nil` value of the second expression is regarded as the absence of a lower boundary, while a `nil` value of the third expression is regarded as the absence of an upper boundary.

### Examples:

```
between(3, 1, 3)           → true
between(3, 1, 2)           → false
```

**implies:** The `implies` function compares two expressions and returns:

- True if the first is false *or* the second is true
- False if the first is true *and* the second is false

### Syntax

#### Input Type

`Boolean implies Boolean`

#### Return Type

Boolean

### Handling Nil Values

The `implies` function returns `nil` if either expression is `nil`.

### Examples:

```
true implies true         → true
true implies false        → false
nil implies false         → nil
```

**in:** The `in` function checks whether a specified search value is in a specified set of values. The return value is `true` if the value is in the set of values and `false` if it is not.

### Syntax

#### Input Type

`Any:1 in(Any:1, Any:1)`

#### Return Type

Boolean

## Handling Nil Values

The function `in` handles `nil` as a normal value.

### Examples:

```
1 in {1,2}           → true
"A" in {"B","C","D","E"} → false
```

**isnil:** The `isnil` function checks whether a value is `nil`.

### Syntax

Input Type	Return Type
<code>isnil(Any:1)</code>	Boolean

## Handling Nil Values

The `isnil` function returns `true` if the value of the expression is `nil` and `false` if not.

### Examples:

```
isnil(nil)           → true
isnil("Hello")       → false
```

**not:** The `not` function returns the negation of a boolean expression.

### Syntax

Input Type	Return Type
<code>not(Boolean)</code>	Boolean

## Handling Nil Values

The `not` function returns `nil` when the Boolean expression is `nil`.

### Examples:

```
not(false)          → true
not(true)           → false
not(nil)            → nil
```

**or:** The `or` function checks whether at least one of two boolean expressions is `true`.

### Syntax

Input Type	Return Type
Boolean or Boolean	Boolean

## Handling Nil Values

If the or expression consists of true and a nil value, the function returns true. If the or expression contains false and a nil value, the result is nil.

### Examples:

true or false	→ true
false or false	→ false
true or true	→ true
true or nil	→ true

## String Functions

This tab contains string functions.

## Handling Nil Values

If one of the expressions has value nil, the functions of the String tab return nil.

### Concatenate: + (concatenate)

The + function concatenates two strings.

### Syntax

Input Type	Return Type
String + Any	String

### Examples:

"child" + "ren"	→ children
"hello " + "anna"	→ hello anna

**chrAt:** The function chrAt returns the character at a specified position of a string. It takes two arguments: the string itself, and an integer representing the position of the character to return. The first character of the string is position 0. Using non-existent positions (smaller than zero or greater than the length minus one) returns a nil value

*DEPRECATED:* Non-ASCII characters become nil. Use: string(s, index, 1).

### Syntax

Input Type	Return Type
chrAt(String, Integer)	String

### Examples:

chrAt("hello world", 0)	→ "h"
-------------------------	-------



`chrAt ("hello world", 1)` → "e"  
`chrAt ("hello world", -1)` → nil

**endsWith:** The `endsWith` function checks whether a string ends with a specified value. It takes two arguments: the string to check, and the value to compare against the string.

### Syntax

Input Type	Return Type
<code>endsWith(String, String)</code>	Boolean

### Examples:

`endsWith("hello", "lo")` → true  
`endsWith("hello", "la")` → false  
`endsWith("hello", "o")` → true

**length:** The `length` function returns an integer representing the length of a string.

### Syntax

Input Type	Return Type
<code>length(String)</code>	Integer

### Examples:

`length("hello")` → 5  
`length("")` → 0

**ltrim:** The `ltrim` function strips white space from the start of a string.

### Syntax

Input Type	Return Type
<code>ltrim(String)</code>	String

### Example:

`ltrim(" hello")` → hello

**rtrim:** The `rtrim` function strips white space from the end of a string.

## Syntax

### Input Type

rtrim(String)

### Return Type

String

### Example:

```
rtrim("hello ")           → hello
```

**search:** The search function checks whether a string contains a specified value. It takes two arguments: the string itself, and the search value to run against the string. The return value is an integer containing the position of the first search value from the beginning of the string. If the string does not contain the specified value, the return value is -1.

The first character of the string is in position 0.

## Syntax

### Input Type

search(String, String)

search(String, Character)

### Return Type

Integer

Integer

### Examples:

```
search("hello, hello", "ll")           → 2
```

```
search("hello, hello", "l")           → 2
```

```
search("hello, hello", "s")           → -1
```

**startsWith:** The startswith function checks whether a string starts with a specified value. It takes two arguments: the string itself, and the value to compare against the string.

## Syntax

### Input Type

startsWith(String, String)

### Return Type

Boolean

### Examples:

```
startsWith("hello", "he")             → true
```

```
startsWith("hello", "ha")             → false
```

```
startsWith("hello", "h")              → true
```

**string:** The string function returns a substring containing characters from a string, starting at a specified position. It takes up to three arguments: the string itself, an integer representing the starting position of the substring, and optionally, an integer specifying the number of characters in the substring.

The first character of the string is in position 0. If the starting position has a negative value, the position is counted from the end of the string instead of the beginning.

### Syntax

Input Type	Return Type
<code>string(String, Integer, Integer)</code>	String
<code>string(String, Integer)</code>	String

### Examples:

```
string("hello", 1, 3)           → "ell"
string("hello", 1)             → "ello"
string("hello", -2)           → "lo"
```

**toLowerCase:** The `toLowerCase` function changes a string to lowercase.

### Syntax

Input Type	Return Type
<code>toLowerCase(String)</code>	String

### Example:

```
toLowerCase("HELLO")           → "hello"
```

**toUpperCase:** The `toUpperCase` function changes a string to uppercase.

### Syntax

Input Type	Return Type
<code>toUpperCase(String)</code>	String

### Example:

```
toUpperCase("hello")           → "HELLO"
```

---

## Importing and exporting IBM SPSS Collaboration and Deployment Services Enterprise View objects

The process of importing and exporting IBM SPSS Collaboration and Deployment Services Enterprise View objects is similar to importing and exporting other repository objects. However, there are a few important considerations:

## Enterprise View

- An imported Enterprise View is considered the master Enterprise View. This means that the imported Enterprise View's LATEST version will be the LATEST in the newly merged Enterprise View (imported LATEST version takes precedence over existing LATEST version). Any matching labels are retained with the imported Enterprise View versions. When viewing the Enterprise View properties, there can be multiple first versions (for example, when the imported and existing views do not share the same ID number). In cases of identical versions, the object timestamp can be used to differentiate between versions.
- The Enterprise View must be exported or imported before importing other IBM SPSS Collaboration and Deployment Services Enterprise View objects. IBM SPSS Collaboration and Deployment Services Enterprise View objects have dependencies on the Enterprise View, requiring that the appropriate Enterprise View is in place before importing.
- The Enterprise View and resource definitions must be exported/imported manually. Exporting the entire content repository will include the resource definitions and Enterprise View, but the import process will not include the resource definitions or Enterprise View.

## Data source definitions

When importing other repository objects, matching data source definitions must exist in the importing repository. The importing repository must contain the same data source definitions that exist in the repository from which the objects were exported. Importing IBM SPSS Collaboration and Deployment Services Enterprise View objects (Enterprise View, Application View, Data Provider Definition, Data Provider Definition - Real Time) does not require that matching data source definitions exist in the importing repository.

## Versions and labels

Although the LATEST label is applied to the LATEST version of the imported Enterprise View, if that version is changed, the new version number is applied to the importing repository Enterprise View.

## Non-Enterprise View users

Non-Enterprise View user profiles can export objects (if they have the export action defined in role) but cannot import objects.

Refer to "Overview" on page 63 for general import/export information.

---

## IBM SPSS Collaboration and Deployment Services Enterprise View data types

The IBM SPSS Collaboration and Deployment Services Enterprise View data types are an approximation of the SQL types defined by the ODBC and JDBC standards. There are a few data type mapping issues you should consider when defining a Data Provider Definition.

Data Provider Definition mappings will not allow a loss of precision. For example, you can map SQL\_INTEGER to the IBM SPSS Collaboration and Deployment Services Enterprise View data type Double, but you cannot map SQL\_DOUBLE to the IBM SPSS Collaboration and Deployment Services Enterprise View data type Integer since the number's decimal portion would be lost.

The exceptions to the no loss of precision rule are the data types Date, Time, and Timestamp. Not all database vendors support separate data types for Date and Time. For instance Microsoft SQL Server does not support any separation of date from time and therefore only supports the data type Datetime (which is exposed as SQL\_TIMESTAMP). Because of this, IBM SPSS Collaboration and Deployment Services Enterprise View allows mapping SQL\_TIMESTAMP fields to the Date, Time, and Timestamp fields. The preferred mapping for SQL\_TIMESTAMP is the IBM SPSS Collaboration and Deployment Services Enterprise

View data type Timestamp. You must set the correct IBM SPSS Collaboration and Deployment Services Enterprise View data type for columns imported from databases that do not separate date from time.

The following tables describes which SQL data types can be mapped to IBM SPSS Collaboration and Deployment Services Enterprise View data types. These mappings are exposed through the “Data Provider Definition editor” on page 218.

Each SQL data type or subtype has one preferred mapping. In the case of SQL\_DECIMAL and SQL\_NUMERIC, the precision and scale are also considered (these data types have several subtypes, each with its own preferred mapping).

*Table 24. Allowed string data type mappings.*

SQL data type	Preferred mapping?
SQL_CHAR	Y
SQL_VARCHAR	Y
SQL_LONGVARCHAR	Y
SQL_WCHAR	Y
SQL_WVARCHAR	Y
SQL_WLONGVARCHAR	Y

*Table 25. Allowed integer data type mappings.*

SQL data type	Preferred mapping?
SQL_INTEGER	Y
SQL_SMALLINT	Y
SQL_TINYINT	Y
SQL_DECIMAL(1-9, 0)	Y
SQL_DECIMAL(1 <= p-s <= 9, s < 0)	Y
SQL_NUMERIC(1-9, 0)	Y
SQL_NUMERIC(1 <= p-s <= 9, s < 0)	Y

*Table 26. Allowed long data type mappings.*

SQL data type	Preferred mapping?
SQL_BIGINT	Y
SQL_INTEGER	N
SQL_SMALLINT	N
SQL_TINYINT	N
SQL_DECIMAL(1-9, 0)	N
SQL_DECIMAL(10-18, 0)	Y
SQL_DECIMAL(1 <= p-s <= 9, s < 0)	N
SQL_DECIMAL(10 <= p-s <= 18, s < 0)	Y
SQL_NUMERIC(1-9, 0)	N
SQL_NUMERIC(10-18, 0)	Y
SQL_NUMERIC(1 <= p-s <= 9, s < 0)	N
SQL_NUMERIC(10 <= p-s <= 18, s < 0)	Y

Table 27. Allowed double data type mappings.

SQL data type	Preferred mapping?
SQL_DOUBLE	Y
SQL_FLOAT	Y
SQL_REAL	N
SQL_INTEGER	N
SQL_SMALLINT	N
SQL_TINYINT	N
SQL_DECIMAL(1-15, 0)	N
SQL_DECIMAL(1 <= p-s <= 15, s < 0)	N
SQL_DECIMAL(1-6, 1-38)	N
SQL_DECIMAL(1-6, 39-308)	Y
SQL_DECIMAL(7-15, 1-308)	Y
SQL_NUMERIC(1-15, 0)	N
SQL_NUMERIC(1 <= p-s <= 15, s < 0)	N
SQL_NUMERIC(1-6, 1-38)	N
SQL_NUMERIC(1-6, 39-308)	Y
SQL_NUMERIC(7-15, 1-308)	Y

Table 28. Allowed float data type mappings.

SQL data type	Preferred mapping?
SQL_REAL	Y
SQL_SMALLINT	N
SQL_TINYINT	N
SQL_DECIMAL(1-6, 0)	N
SQL_DECIMAL(1-6, 1-38)	Y
SQL_DECIMAL(1 <= p-s <= 6, s < 0)	N
SQL_NUMERIC(1-6, 0)	N
SQL_NUMERIC(1-6, 1-38)	Y
SQL_NUMERIC(1 <= p-s <= 6, s < 0)	N

Table 29. Allowed decimal data type mappings.

SQL data type	Preferred mapping?
SQL_DECIMAL(16+, 1+)	Y
SQL_DECIMAL(19+, 0)	Y
SQL_DECIMAL(*, 309+)	Y
SQL_DECIMAL(1-15, 0-308)	N
SQL_DECIMAL(1 <= p-s <= 15, s < 0)	N
SQL_NUMERIC(16+, 1+)	Y
SQL_NUMERIC(19+, 0)	Y
SQL_NUMERIC(*, 309+)	Y
SQL_NUMERIC(1-15, 0-308)	N

Table 29. Allowed decimal data type mappings (continued).

SQL data type	Preferred mapping?
SQL_NUMERIC(1 <= p-s <= 15, s < 0)	N
SQL_INTEGER	N
SQL_SMALLINT	N
SQL_TINYINT	N
SQL_BIGINT	N
SQL_DOUBLE	N
SQL_FLOAT	N
SQL_REAL	N

Table 30. Allowed boolean data type mappings.

SQL data type	Preferred mapping?
SQL_BIT	Y

Table 31. Allowed date data type mappings.

SQL data type	Preferred mapping?
SQL_DATE	Y
SQL_TIMESTAMP	N

Table 32. Allowed time data type mappings.

SQL data type	Preferred mapping?
SQL_TIME	Y
SQL_TIMESTAMP	N

Table 33. Allowed timestamp data type mappings.

SQL data type	Preferred mapping?
SQL_TIMESTAMP	Y

Note: IBM SPSS Collaboration and Deployment Services Enterprise View does not support SQL data types not specifically listed in the preceding table.

Many IBM SPSS Collaboration and Deployment Services Enterprise View data types allow you to map multiple SQL data types. The SQL data type provided to the client application may change depending on which Data Provider Definition is in use. Client applications must dynamically handle data type changes and should not make assumptions about data types that are provided at runtime.

When you import physical table column definitions into the Enterprise View (through the “Adding columns to the Enterprise View” on page 206 process), a given SQL data type will map to only one IBM SPSS Collaboration and Deployment Services Enterprise View data type. This mapping is determined by the SQL data type's preferred mapping.

Table 34. Preferred SQL data type mappings.

SQL data type	Preferred IBM SPSS Collaboration and Deployment Services Enterprise View data type
SQL_CHAR	String

Table 34. Preferred SQL data type mappings (continued).

SQL data type	Preferred IBM SPSS Collaboration and Deployment Services Enterprise View data type
SQL_VARCHAR	String
SQL_LONGVARCHAR	String
SQL_WCHAR	String
SQL_WVARCHAR	String
SQL_WLONGVARCHAR	String
SQL_INTEGER	Integer
SQL_SMALLINT	Integer
SQL_TINYINT	Integer
SQL_BIGINT	Long
SQL_DOUBLE	Double
SQL_FLOAT	Double
SQL_REAL	Float
SQL_DECIMAL(1-9, 0)	Integer
SQL_DECIMAL(1 <= p-s <= 9, s < 0)	Integer
SQL_DECIMAL(10-18, 0)	Long
SQL_DECIMAL(10 <= p-s <= 18, s < 0)	Long
SQL_DECIMAL(1-6, 1-38)	Float
SQL_DECIMAL(1-6, 39-308)	Double
SQL_DECIMAL(7-15, 1-308)	Double
SQL_DECIMAL(16+, * )	Decimal
SQL_DECIMAL(*, 309+)	Decimal
SQL_NUMERIC(1-9, 0)	Integer
SQL_NUMERIC (1 <= p-s <= 9, s < 0)	Integer
SQL_NUMERIC (10-18, 0)	Long
SQL_NUMERIC (10 <= p-s <= 18, s < 0)	Long
SQL_NUMERIC (1-6, 1-38)	Float
SQL_NUMERIC (1-6, 39-308)	Double
SQL_NUMERIC (7-15, 1-308)	Double
SQL_NUMERIC (16+, * )	Decimal
SQL_NUMERIC (*, 309+)	Decimal
SQL_BIT	Boolean
SQL_DATE	Date
SQL_TIME	Time
SQL_TIMESTAMP	Timestamp

Although a SQL data type may have a corresponding IBM SPSS Collaboration and Deployment Services Enterprise View data type, the preferred mapping may not always be to the same IBM SPSS Collaboration and Deployment Services Enterprise View data type. Precision and scale are also considered. This is due to the fact that vendors such as Oracle expose all of their non-floating point



numbers as SQL\_DECIMAL. Basing the preferred mappings at least partially on precision and scale is necessary to ensure that all IBM SPSS Collaboration and Deployment Services Enterprise View data types can be imported from physical tables.

*Note:* The preferred mapping for SQL\_DECIMAL is whichever IBM SPSS Collaboration and Deployment Services Enterprise View data type best matches the scale and precision. Only in cases where there is no appropriate match to another data type will the preferred mapping for SQL\_DECIMAL map to the data type Decimal. However, this does not mean that you cannot map any SQL\_DECIMAL to the IBM SPSS Collaboration and Deployment Services Enterprise View data type Decimal. On UNIX-based operating systems, SQL\_BIGINT type is treated as numeric(19,0), which may result in loss of precision.



---

## 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 want 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.

---

## Privacy policy considerations

IBM Software products, including software as a service solutions, (“Software Offerings”) may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering’s use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM’s Privacy Policy at <http://www.ibm.com/privacy> and IBM’s Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled “Cookies, Web Beacons and Other Technologies” and the “IBM Software Products and Software-as-a-Service Privacy Statement” at <http://www.ibm.com/software/info/product-privacy>.

---

## Trademarks

IBM, the IBM logo, and [ibm.com](http://www.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](http://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.



---

## Glossary

This glossary includes terms and definitions for IBM SPSS Collaboration and Deployment Services.

The following cross-references are used in this glossary:

- See refers you from a term to a preferred synonym, or from an acronym or abbreviation to the defined full form.
- See also refers you to a related or contrasting term.

To view glossaries for other IBM products, go to [www.ibm.com/software/globalization/terminology](http://www.ibm.com/software/globalization/terminology) (opens in new window).

---

### A

#### access control list (ACL)

In computer security, a list associated with an object that identifies all the subjects that can access the object and their access rights.

**ACL** See access control list.

**action** A permission for an aspect of system functionality. For example, the ability to set up notifications is defined as an action. Actions are grouped and assigned to users through roles. See also role.

#### Active Directory (AD)

A hierarchical directory service that enables centralized, secure management of an entire network, which is a central component of the Microsoft Windows platform.

**AD** See Active Directory.

#### allowed user

A subset of the users defined in a remote directory, such as SiteMinder or Windows Active Directory, that are allowed access to SPSS Predictive Enterprise Services. Allowed users are defined when only a few users in a remote directory need access to the application.

**API** See application programming interface.

#### appender

A component that receives logging

requests from a logger and writes log statements to a specified file or console. See also logger.

#### application programming interface (API)

An interface that allows an application program that is written in a high-level language to use specific data or functions of the operating system or another program.

#### application view

A logical mapping of data items defined in the enterprise view to a specific application. It captures additional information about how tables and columns are used in the context of an application (for example, call center or fraud detection). There can be more than one application view associated with the enterprise view.

---

### B

#### batch file

A file that contains instructions that are processed sequentially, as a unit.

#### binary large object (BLOB)

A data type whose value is a sequence of bytes that can range in size from 0 bytes to 2 gigabytes less 1 byte. This sequence does not have an associated code page and character set. BLOBs can contain, for example, image, audio, or video data.

**BLOB** See binary large object.

#### break group

A set of rows of returned data that are grouped according to a common column value. For example, in a column of states, the rows of data for each state are grouped together.

#### burst report

A report that generates multiple output files during a single run by using multiple input parameters taken from break groups in the report.

---

## C

### **cascading permission**

A permission of a parent folder in the content repository that has been propagated to its child objects.

### **character large object (CLOB)**

A data type whose value is a sequence of characters (single byte, multibyte, or both) that can range in size from 0 bytes to 2 gigabytes less 1 byte. In general, the CLOB data type is used whenever a character string might exceed the limits of the VARCHAR data type.

**CLOB** See character large object.

### **common warehouse metamodel (CWM)**

A metamodel written to be a common standard by the Object Management Group (OMG).

### **content repository**

A centralized location for storing analytical assets, such as models and data. Content repository includes facilities for security and access control, content management, and process automation.

### **context data**

Input data that is passed with a scoring request in real time. For example, when a score is requested for a customer based on credit rating and geocode, the credit score and geocode will be the context data for the request.

### **credential**

Information acquired during authentication that describes a user, group associations, or other security-related identity attributes, and that is used to perform services such as authorization, auditing, or delegation. For example, a user ID and password are credentials that allow access to network and system resources.

**CWM** See common warehouse metamodel.

---

## D

### **Data Provider Definition**

A mapping of the logical column definitions in the application view to physical table columns in the data source. A Data Provider Definition links source data directly to an application view and

indirectly to enterprise view tables and columns. A single application view table may have more than one Data Provider Definition, each of which may point to multiple data sources.

### **data warehouse**

A subject-oriented collection of data that is used to support strategic decision making. The warehouse is the central point of data integration for business intelligence. It is the source of data for data marts within an enterprise and delivers a common view of enterprise data.

### **distinguished name (DN)**

The name that uniquely identifies an entry in a directory. A distinguished name is made up of attribute:value pairs, separated by commas. For example, CN=person name and C=country or region.

**DN** See distinguished name.

### **Document Object Model (DOM)**

A system in which a structured document, for example an XML file, is viewed as a tree of objects that can be programmatically accessed and updated. See also Simple API for XML.

### **document type definition (DTD)**

The rules that specify the structure for a particular class of SGML or XML documents. The DTD defines the structure with elements, attributes, and notations, and it establishes constraints for how each element, attribute, and notation can be used within the particular class of documents.

**DOM** See Document Object Model.

### **dormant schedule**

A schedule associated with a deleted or unlabeled version of a job. A dormant schedule cannot be used until it is associated with a valid labeled job version.

**DTD** See document type definition.

---

## E

**EAR** See enterprise archive.

### **enterprise archive (EAR)**

A specialized type of JAR file, defined by



the Java EE standard, used to deploy Java EE applications to Java EE application servers. An EAR file contains EJB components, a deployment descriptor, and web archive (WAR) files for individual web applications. See also Java archive, web archive.

**enterprise view**

An abstracted dictionary of enterprise data. The enterprise view provides a simplified business-oriented picture of the data in different data sources across the enterprise for data mining and reporting. Enterprise view is populated with data items defined as logical tables and columns. Enterprise view data items are mapped to analytical applications through the application views and to physical data sources through Data Provider Definitions.

**execution server**

A server that enables analytical processing of resources stored in the repository. For example, to execute an IBM SPSS Statistics syntax in an IBM SPSS Collaboration and Deployment Services job, an IBM SPSS Statistics execution server must be designated.

**export** The process of storing objects and metadata from the content repository to an external file.

**extended group**

A locally-defined group of remote users. Extended groups are defined when groups in the remote directory are not fine-grained enough.

**Extensible Markup Language (XML)**

A standard metalanguage for defining markup languages that is based on Standard Generalized Markup Language (SGML).

**Extensible Stylesheet Language (XSL)**

A language for specifying style sheets for XML documents. Extensible Stylesheet Language Transformation (XSLT) is used with XSL to describe how an XML document is transformed into another document.

---

**F**

**field content assist**

A feature that provides predefined system and variable values for entry fields.

---

**G**

**general job step**

A method for running native operating system commands and executable programs on a host or a remote process server. General jobs have access to files stored within the repository and on the file system and can be used to control the input/output of analytical processing.

---

**I**

**import**

The process of adding objects and metadata defined in an external file generated by export, to the content repository.

**iterative consumer reporting job step**

A job step that is passed a set of input values generated by a preceding iterative producer reporting job step. The report in iterative consumer job step is executed for each tuple in the received data set.

**iterative producer reporting job step**

A job step that generates a set of values passed as input parameters to a following iterative consumer job step.

---

**J**

**JAAS** See Java Authentication and Authorization Service.

**JAR** See Java archive.

**Java archive (JAR)**

A compressed file format for storing all of the resources that are required to install and run a Java program in a single file. See also enterprise archive, web archive.

**Java Authentication and Authorization Service (JAAS)**

In Java EE technology, a standard API for performing security-based operations. Through JAAS, services can authenticate and authorize users while enabling the applications to remain independent from underlying technologies.

### **Java Generic Security Services (JGSS)**

A specification that provides Java programs access to the services that include the signing and sealing of messages and a generic authentication mechanism.

### **Java Naming and Directory Interface (JNDI)**

An extension to the Java platform that provides a standard interface for heterogeneous naming and directory services.

**JGSS** See Java Generic Security Services.

**JNDI** See Java Naming and Directory Interface.

**job** A mechanism for automating analytical processing. A job consists of job steps, executed sequentially or conditionally. Input parameters can be defined for a job. A job can be run on demand or triggered by time-based or message-based schedules, with records of job execution stored as job history.

### **job step**

A discrete unit of processing in a job. Depending on the type, job steps can be run on the content repository host or specially defined execution or remote process servers. Objects stored in the repository or the file system can provide input for a job step, and job step output can be stored in the repository or written to the file system.

---

## **K**

**KDC** See key distribution center.

### **Kerberos**

A network authentication protocol that is based on symmetric key cryptography. Kerberos assigns a unique key, called a ticket, to each user who logs on to the network. The ticket is embedded in messages that are sent over the network. The receiver of a message uses the ticket to authenticate the sender.

### **key distribution center (KDC)**

A network service that provides tickets and temporary session keys. The KDC maintains a database of principals (users and services) and their associated secret keys. It is composed of the authentication server and the ticket granting ticket server.

### **keystore**

In security, a file or a hardware cryptographic card where identities and private keys are stored, for authentication and encryption purposes. Some keystores also contain trusted or public keys.

---

## **L**

**LDAP** See Lightweight Directory Access Protocol.

### **Lightweight Directory Access Protocol (LDAP)**

An open protocol that uses TCP/IP to provide access to directories that support an X.500 model and that does not incur the resource requirements of the more complex X.500 Directory Access Protocol (DAP). For example, LDAP can be used to locate people, organizations, and other resources in an Internet or intranet directory.

**lock** The process by which integrity of data is ensured by preventing more than one user from accessing or changing the same data or object at the same time.

**logger** A component that prepares log statements to be written to console or log file. See also appender.

---

## **M**

### **message-based schedule**

A schedule that is used to trigger job execution by an event signalled by a Java Messaging Service (JMS) message. For example, when a job relies on the input from a third-party application, the application must send a JMS message when the input file is ready for processing.

### **metamodel**

A model that defines the language for expressing a model.

### **meta-object**

An instance of an XMI class as defined in the metamodel.

### **meta-object facility (MOF)**

A generalized facility and repository for storing abstract information about concrete object systems; dealing mostly with construction, standardized by the Object Management Group (OMG).

**MIME** See Multipurpose Internet Mail Extensions.

**MOF** See meta-object facility.

**Multipurpose Internet Mail Extensions (MIME)**  
An Internet standard that allows different forms of data, including video, audio, or binary data, to be attached to email without requiring translation into ASCII text.

---

## N

**notification**  
A mechanism that is used to generate email messages informing users of specific types of system events, such as changes to content repository objects and processing success and failure. Unlike subscriptions, notifications can be set up to send email to multiple users.

---

## O

**Object Management Group (OMG)**  
A non-profit consortium whose purpose is to promote object-oriented technology and the standardization of that technology. The Object Management Group was formed to help reduce the complexity, lower the costs, and hasten the introduction of new software applications.

**ODS** See Output Delivery System.

**OMG** See Object Management Group.

**Output Delivery System (ODS)**  
A method of controlling the destination for output within SAS. ODS can route SAS output to a SAS data file, a text listing file, HTML files, and files optimized for high-resolution printing.

---

## P

**package**  
An installable unit of a software product. Software product packages are separately installable units that can operate independently from other packages of that software product.

**principal**  
An entity that can communicate securely with another entity. A principal is

identified by its associated security context, which defines its access rights.

---

## R

**remote process server**  
A remote system that is designated for running native operating system commands and executable programs.

**repository content adapter**  
An optional software package that enables storing and processing content from other IBM SPSS applications, such as Statistics, Modeler, and Data Collection, as well as third parties.

**repository database**  
A relational database that is used for storing content repository objects and metadata.

**resource**  
A content repository object.

**resource definition**  
A subset of content repository resources used to enable analytical processing, such as definitions of data sources, credentials, execution servers, and JMS message domains.

**role** A set of permissions or access rights. See also action.

---

## S

**SAX** See Simple API for XML.

**schedule**  
A content repository object that triggers job execution.

**scoring configuration**  
A configuration that defines model-specific settings for generating real-time scores, such as input data, processing rules, outputs, logging, etc.

**security provider**  
A system that performs user authentication. 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.

**service provider interface (SPI)**  
An API that supports replaceable

components and can be implemented or extended by a third party.

## **SGML**

See Standard Generalized Markup Language.

## **shell script**

A program, or script, that is interpreted by the shell of an operating system.

## **Simple API for XML (SAX)**

An event-driven, serial-access protocol for accessing XML documents, used. A Java-only API, SAX is used by most servlets and network programs to transmit and receive XML documents. See also Document Object Model.

## **single sign-on (SSO)**

An authentication process in which a user can access more than one system or application by entering a single user ID and password.

**SOAP** A lightweight, XML-based protocol for exchanging information in a decentralized, distributed environment. SOAP can be used to query and return information and invoke services across the Internet.

**SPI** See service provider interface.

**SSO** See single sign-on.

## **Standard Generalized Markup Language (SGML)**

A standard metalanguage for defining markup languages that is based on the ISO 8879 standard. SGML focuses on structuring information rather than presenting information; it separates the structure and content from the presentation. It also facilitates the interchange of documents across an electronic medium.

## **stop word**

A word that is commonly used, such as "the," "an," or "and," that is ignored by a search application.

## **subscription**

Email notices and Really Simple Syndication (RSS) feeds that repository users create to receive when the state of an asset changes.

---

## **T**

**TGT** See ticket-granting ticket.

## **ticket-granting ticket (TGT)**

A ticket that allows access to the ticket granting service on the key distribution center (KDC). Ticket granting tickets are passed to the principal by the KDC after the principal has completed a successful request. In a Windows 2000 environment, a user logs on to the network and the KDC will verify the principal's name and encrypted password and then send a ticket granting ticket to the user.

## **time-based schedule**

A schedule that triggers job execution at a specified time or date. For example, a time-based schedule may run a job at 5:00 pm every Thursday.

---

## **U**

## **Universally Unique Identifier (UUID)**

The 128-bit numeric identifier that is used to ensure that two components do not have the same identifier.

**UUID** See Universally Unique Identifier.

---

## **V**

## **Velocity**

A Java-based template engine that provides a simple and powerful template language to reference objects defined in Java code. Velocity is an open source package directed by the Apache Project.

---

## **W**

**W3C** See World Wide Web Consortium.

**WAR** See web archive.

## **web archive (WAR)**

A compressed file format, defined by the Java EE standard, for storing all the resources required to install and run a web application in a single file. See also enterprise archive, Java archive.

## **Web Services Description Language (WSDL)**

An XML-based specification for describing networked services as a set of endpoints operating on messages

containing either document-oriented or procedure-oriented information.

**World Wide Web Consortium (W3C)**

An international industry consortium set up to develop common protocols to promote evolution and interoperability of the World Wide Web.

**WSDL**

See Web Services Description Language.

---

**X**

**XMI** See XML Metadata Interchange.

**XML** See Extensible Markup Language.

**XML Metadata Interchange (XMI)**

A model-driven XML integration framework for defining, interchanging, manipulating, and integrating XML data and objects. XMI-based standards are in use for integrating tools, repositories, applications, and data warehouses.

**XSL** See Extensible Stylesheet Language.



---

# Index

## A

- A/B split testing
  - for scoring 111, 112
- abs function 236
- accessibility 199
  - blind users 201
  - keyboard navigation 199
  - visually impaired users 201
- accessing
  - expired files 33
- actions
  - promotion 69
- add data set
  - Application Server Data Source 226
  - connection settings 225
  - Data Service data source 226
  - JDBC data source 225
  - select tables to add 225
- adddays function 240
- adding
  - job variables 121
  - principals for label 34
  - promotion policies 55
- adding columns
  - connection settings 205
  - select columns 205
- adding steps 122
- adding tables
  - connection settings 205
  - select columns 205
- addition function 235
- addmonths function 240
- addyears function 240
- advanced mode
  - for analytic data views 84
- advanced search 17, 18, 19, 20, 21
- aggregate functions 228, 231
  - handling nil values 229
- aliases
  - creating 106
  - deleting 107
  - editing 106
  - for scoring configurations 105
  - in A/B split testing 111, 112
  - viewing 113
- analytic data view
  - champion challenger 195
- analytic data views 71, 78
  - adding attributes 85
  - adding tables 84
  - advanced mode 84
  - business object models 90, 91, 92
  - creating 72
  - data access plans 73, 75, 78, 79, 81
  - data mapping 77
  - data models 75, 82, 88, 90
  - deleting attributes 90
  - deleting tables 90
  - derived attributes 87
  - execution object models 93
  - table relationships 86

- and function 247
- any data type 229
- any:1 data type 229
- application
  - exiting 11
  - navigating 9
  - starting 9
- Application Server Data Source
  - add data set 226
- application server data sources 49, 52
  - for data access plans 79
- Application View 212
  - changing 219
  - column attributes 213
  - column count 213
  - columns 212
  - create 212
  - merging versions 216, 217
  - removing columns 212
  - removing tables 212
  - saving 215
  - tables 212
- Application View editor 212
- applying
  - locks 22
- arithmetic functions 234
- association sets
  - creating 110
  - deleting 111
  - editing 110
  - for scoring configurations 107
  - in A/B split testing 111, 112
  - viewing 113
- ATOM feeds 36
- attachment 149, 150, 153
- auditing
  - for scoring 102
- avg function 232

## B

- basic functions 228
  - handling nil values 229
- batch
  - data access plans 75
- batch files 173
- batch scoring 104
  - between function 248
- BIRT Report Designer for IBM SPSS 2, 6
- BIRT steps 157
  - adding to jobs 157
  - burst 158
  - clean-up 160
  - data source 157
  - defining general properties 157
  - iterative 158
  - locale 159
  - metadata 159
  - naming 157
  - notifications 161
  - output file format 159

- BIRT steps (*continued*)
  - output file location 159
  - output file name 159
  - parameters 159
  - prompts 159
  - rendering options 159
  - result permissions 159
  - results 159
  - single 158
  - type 158
  - variables 159
  - versions 157
- bit function 237
- blank cells
  - in job history 137
- BOM
  - See* business object models
- bulk property updates 45
  - general 45
  - permission 46
  - version 45
- business object models 90
  - exporting 84, 91
  - importing 84, 92

## C

- cachable 222
- cache sizes
  - for scoring 104
- cache timeout 222
- canceling
  - jobs 137
- cardinality
  - data model relationships 82
- categories
  - column attributes 208
  - exporting 209
  - removing 209
- cells
  - blank 137
- challenger models 189, 190, 191, 192
  - data sources 192
- champion challenger 138, 139, 190, 191, 192, 193, 194, 195, 196
- champion models 189, 190, 192, 193
- changing
  - Application View 219
  - IBM SPSS Collaboration and Deployment Services Enterprise View 216
  - passwords 15
- chrAt function 250
- clean up 160, 165
- client
  - starting 9
- clusters 60
- Cognos
  - champion challenger 196
- collaboration 1

- column
  - select columns 205
- column attributes
  - Application View 213
  - categories 208
  - description 209, 215
  - IBM SPSS Collaboration and Deployment Services Enterprise View 207
  - min/max 209
- column count 213
- columns 53
  - Application View 212
  - IBM SPSS Collaboration and Deployment Services Enterprise View 206
- comparing
  - models 189, 190
- computation time
  - for scoring configurations 102
- computation wait time
  - for scoring configurations 102
- concatenate function 250
- conditional connector 124
- configuration objects 218
- configurations
  - for scoring models 98
- conflict resolution 66
  - duplicates 67, 68
  - individual 67, 68
  - invalid versions 67, 68
- conflicts
  - duplicates 67, 68
  - import 66, 67, 68
  - invalid versions 67, 68
- connection
  - existing 15
  - new 14
  - server 14, 15
  - terminating 15
- connection settings
  - add data set 225
  - adding columns 205
  - adding tables 205
  - select data set to add 226
- connections
  - server 33, 37
- connectors
  - conditional 124
  - deleting 126
  - fail 124
  - pass 124
  - relationship 123, 124, 125, 126
  - sequential 124
- constant filters
  - for data access plans 79
- content assist 10
- content explorer
  - deleting files 17
  - files 16, 17
  - jobs 119, 120
  - keyboard navigation 199
  - organization 13
  - overview 13
  - permissions 16
  - properties 25
  - search 17, 18, 19, 20, 21

- content explorer (*continued*)
  - searching 21
  - server connections 14, 15
  - versions 30, 31, 37
- content object
  - properties 25
- content objects
  - definition 13
  - permissions 16
  - storage 13, 16, 17
  - versions 29, 31
- content repository 13, 16, 17, 197, 198
  - deleting files 17
- Content Repository
  - execution server 58
  - server definition 58
- context data sources
  - for data access plans 79
- conventions
  - naming 10
- conversion functions 237
- copying
  - topics 43
- count function 232
- create
  - Application View 212
  - Data Provider Definition 218
  - Data Provider Definition - Real Time 221
- creating
  - custom properties 38, 39
  - schedules 130
  - scoring configuration aliases 106
  - scoring configuration associations 110
  - server clusters 60
  - topics 42, 43
- credential definitions 47, 48, 204
- credential destination 47
- credentials 47, 55
  - for schedules 130
  - job 118
  - promoting 69
- custom properties
  - accessing 38
  - creating 27, 38, 39, 40
  - defining 27
  - deleting 41
  - editing 27, 40
  - existing 27, 40
  - label 39
  - modifying 27, 40
  - new 38
  - overview 38
  - property type 39
  - property types 27
  - searching 40
  - selection values 27, 39, 40
  - setting values 38
- customizing
  - notification messages 153

## D

- daily schedules 131

- data
  - filtering model management views 144
- data access plans 72, 73
  - creating 75
  - deleting 81
  - for analytic data views 75, 78, 79, 81
  - for scoring 99
  - overriding data sources 78, 79
  - previewing data 81
  - real-time 79
- data access time
  - for scoring configurations 102
- data files 118
  - champion challenger 194
- data initialization time
  - for scoring configurations 102
- data models 72
  - adding attributes 85
  - adding derived attributes 87
  - adding tables 75, 84
  - data mapping 77
  - defining relationships 86
  - deleting attributes 90
  - deleting tables 90
  - for analytic data views 82
  - modifying properties 88
  - relationships 82
  - stream fields 78
  - table attributes 78
- Data Provider Definition 157, 217
  - configuration objects 218
  - create 218
  - mapping columns 219
  - saving 221
  - select Application View 218
  - validating 220
- Data Provider Definition - Real Time 221
  - create 221
  - defining data sets 225
  - editing data sets 226
  - for scoring 99
  - mapping data set columns 222
  - preview data 223
  - saving 227
  - select Application View 222
  - working with data sets 224
- Data Provider Definition - Real Time editor 222
- Data Provider Definition editor 218
- Data Service data source
  - add data set 226
- data service data sources 49, 52
  - for data access plans 79
  - keys 53
  - tables 53
- data source
  - BIRT 157
  - reporting step 157
- data source definitions 47, 204
- Data Source Definitions 253
- data sources 48
  - application server data sources 49, 52
  - data service data sources 49, 52, 53
  - JDBC data sources 49, 51



- data sources (*continued*)
  - modifying 54
  - ODBC data sources 49
  - promoting 69
- data types 53, 229, 254
- databases 118
- date and time functions 240
- date function 237, 241
- date range 18, 19, 20
- dates
  - expiration 31, 32, 33
- day function 241
- dayofweek function 241
- dayofyear function 242
- dbl function 237
- decimal function 238
- defining data sets
  - Data Provider Definition - Real Time 225
- defining tables 52
- delayed promotion 56, 57
- delete search terms 21
- deleting 21
  - custom properties 41
  - files 17
  - groups 29
  - job variables 122
  - principals for label 34
  - schedules 133
  - scoring configuration aliases 107
  - scoring configuration associations 111
  - scoring configurations 114
  - search terms 21
  - topics 44
  - users 29
  - versions 31
- delivery failure 155
- deployment 2
- description 37
  - column attributes 209, 215
- destination
  - credentials 47
  - server 58
- destination name 54, 55
- DEVICE 168, 169
- diff function 242
- distribution channels 36
- division function 236
- drag and drop 9
- duplicates 67, 68
- durable subscriptions 131
- Dynamic notification lists 152, 153

## E

- edit search terms 20
- editing 20
  - custom properties 27, 40
  - job variables 122
  - permissions for label 34
  - properties 26
  - schedules 133
  - scoring configuration aliases 106
  - scoring configuration associations 110
  - scoring configurations 114

- editing (*continued*)
  - server clusters 61
  - topics 27
- editing data sets
  - Data Provider Definition - Real Time 226
- editor
  - Application View 212
  - Data Provider Definition 218
  - Data Provider Definition - Real Time 222
  - IBM SPSS Collaboration and Deployment Services Enterprise View 204
- editors
  - Application View 203
  - Data Provider Definition 203
  - IBM SPSS Collaboration and Deployment Services Enterprise View 203
- email delivery failure 155
- enabling
  - filtering 142, 143
- end date
  - for schedules 131
- endsWith function 251
- enter key 9
- Enterprise View
  - champion challenger 195
  - merging versions 216, 217
- equal to function 246
- error messages
  - for schedules 136
- evaluation
  - results 138, 139, 140
- evaluation type
  - filtering model management views 144
- exclusions
  - search terms 17, 21
- executables 173
- executing jobs 123
- execution object models 93
  - exporting 93
- execution servers 5
  - remote process 2, 5
  - remote process servers 167, 173
  - SAS 2, 5, 167
- exist function 232
- existing jobs 120
- exiting
  - application 11
  - client 11
  - system 11
- expiration
  - submitted jobs 198
- expiration date 37
- expiration dates 198
  - exporting 32
  - importing 32
  - modifying 32
  - overview 31
  - reactivating 33
  - restrictions 32
  - searching 33
  - setting 32
  - viewing 33

- export 147
  - external references 63
  - folders 63, 64, 66, 67, 68
  - jobs 63
- exporting 63, 64, 66, 67, 68
  - business object models 84, 91
  - categories 209
  - conflicts 66, 67
  - execution object models 93
  - expiration dates 32
  - resource definitions 62
  - restrictions 64
- expression editor 227
  - expression panel 227
  - functions panel 228
- expression panel 227
  - of expression editor 227
- expression status field 229
- external files
  - in SAS steps 169
- external references 63
  - exporting 63
  - importing 63

## F

- F1 help 9
- fail connector 124
- field content assist 10
- file
  - filtering model management views 144
  - permissions 127
  - subscriptions 154
  - versions 29, 30, 31, 37
- FILENAME 169
- fileref 169
- files 16
  - .pes 63, 64
  - adding 16
  - adding to repository 63
  - copying 16
  - deleting 17
  - downloading 17, 63
  - expired 33
  - expiring 31, 32, 33
  - locking 22
  - moving 16
  - opening 16
  - permissions 16
  - reactivating 33
  - unlocking 22
- filtering
  - enabling 142, 143
  - job history 143
  - job schedule 142
  - scoring view 114
- filters 141, 142, 143
  - data 144
  - evaluation type 144
  - file 144
  - index 144
  - job 141
  - label 141
  - model execution date 144
  - trend 144
- flt function 238

- folder
  - resource definitions 47, 48, 57, 58, 59
  - search 17, 18, 19, 20, 21
  - searching 21
- folder options 150
- folders
  - child 28
  - content repository 197, 198
  - expiration of 198
  - exporting 63, 64, 66, 67, 68
  - importing 63, 64, 65, 66, 67, 68
  - parent 28
  - resource definitions 13
  - restrictions 197
  - submitted jobs 197, 198
- functions
  - introduction 228
  - types 228
- functions panel 227, 228
  - of expression editor 228

## G

- General job steps 173
  - adding to jobs 173
  - defining properties 173
  - examples 175, 176
  - input files 174
  - naming 173
  - output files 174, 175
  - working directory 173, 174, 175
- general properties 26
- getting started 203
- global conflict resolution 66
- glossary 265
- greater than function 246
- greater than or equal to function 247
- group search terms 20
- groups 128
  - deleting 29
  - existing 28
  - new 28
  - permissions 127
- GSFNAME 169
- guidelines
  - naming 10

## H

- help
  - accessing 9
  - F1 9
  - keyboard navigation 200
- history
  - job 137, 143
  - job step 138
  - server status 145
- hourly schedules 131
- HTML output
  - for IBM SPSS Statistics steps 169
- HTMLC 159

## I

- IBM Operational Decision Manager
  - business object models 90, 91, 92

- IBM Operational Decision Manager
  - (continued)
  - execution object models 93
- IBM SPSS Collaboration and Deployment Services Deployment Manager 2, 4
- IBM SPSS Collaboration and Deployment Services Deployment Portal 2, 4
- IBM SPSS Collaboration and Deployment Services Enterprise View 2, 5, 204
  - changing 216
  - column attributes 207
  - columns 206
  - credential definitions 204
  - data source definitions 204
  - Data Source Definitions 253
  - data types 254
  - editors 203
  - manage keys 209
  - manage links 210
  - removing columns 207
  - removing tables 206
  - saving 211
  - tables 205
- IBM SPSS Collaboration and Deployment Services Enterprise View editor 204
- IBM SPSS Collaboration and Deployment Services Repository 2, 3
- ifnil function 243
- immediate promotion 56, 57
- implies function 248
- import 147
  - external references 63
  - folders 63, 64, 65, 66, 67, 68
  - jobs 63
  - security permissions 64
- import order 64
- importing 63, 64, 65, 66, 67, 68
  - business object models 84, 92
  - conflicts 66, 67
  - expiration dates 32
  - resource definitions 62
  - restrictions 64
- in function 248
- index
  - filtering model management views 144
- individual conflict resolution 67, 68
- inheritance
  - permissions 28
- input files
  - General job steps 174
- int function 238
- invalid versions 67, 68
- isnl function 249
- iterative principals 128

## J

- JDBC data source
  - add data set 225
- JDBC data sources 49
  - for data access plans 79
  - third-party drivers 51
- JMS 54, 55, 183
  - mapping variables 132
- JMS message domain 133

- job
  - notifications 147
  - subscriptions 154
- Job Editor
  - keyboard navigation 200
- job history
  - filtering 143
  - keyboard navigation 200
  - limiting 143
  - reordering 143
  - viewing 137, 143
- job labels
  - for schedules 130
- job output 138
- job schedule
  - filtering 142
  - keyboard navigation 200
  - limiting 142
  - reordering 142
  - viewing 142
- job step
  - notifications 148
  - permissions 127
- job step history
  - viewing 137, 138
- job step name
  - champion challenger 190
- job step output 127
- job steps
  - history 138
  - results 138
- job variables
  - in schedules 132
  - logging values 138
- jobs
  - adding steps 120, 122
  - canceling 137
  - components 117
  - content explorer 119, 120
  - creating 119
  - definition 117
  - dependencies 118
  - editing 120, 122
  - existing 120
  - exporting 63
  - filtering 142, 143
  - history 137, 138, 143
  - importing 63
  - job editor 120, 122, 123, 124, 125, 126
  - logs 138
  - new 119
  - notifications 129
  - opening 119, 120
  - overview 117
  - prerequisites 118
  - process 118
  - promoting 69
  - properties 120
  - relationships 123, 124, 125, 126
  - running 123, 129, 130, 131, 136
  - saving 126
  - schedule 142
  - scheduling 129, 131, 136
  - searching 142, 143
  - status 137, 138
  - submitted 197, 198
  - variables 121, 129

## K

- keyboard navigation 199
  - in content explorer 199
  - in help system 200
  - in job editor 200
  - in job history view 200
  - in tables 200
- keys 209
- keywords 37

## L

- label
  - notifications 149
- labels
  - applying 30
  - filtering 142, 143
  - for versions 29, 30, 31
  - LATEST 29
  - permissions 34
  - recommendations 30
  - removing 30
  - security 34
  - server version 37
- LATEST label 29
- length function 251
- less than function 245
- less than or equal to function 246
- links 203, 210
- lng function 239
- locking objects 22
- locks
  - applying 22
  - files 22
  - objects 22
  - removing 22
  - resource definitions 22
- log in
  - server 15
- log off
  - server 15
- logging
  - for scoring 102
- logs
  - job 138
  - output 138
- Lotus 151
- ltrim function 251

## M

- manage keys
  - IBM SPSS Collaboration and Deployment Services Enterprise View 209
- manage label permission
  - for labels 34
- manage links
  - IBM SPSS Collaboration and Deployment Services Enterprise View 210
- map columns 219, 222
- mapping columns
  - Data Provider Definition 219

- mapping data set columns
  - Data Provider Definition - Real Time 222
- max function 233
- merging
  - version conflicts 217
  - view versions 216, 217
- message customization 149, 150
  - body text 153
  - errors 153
  - event property variables 153
  - HTML formatting 153
  - message customization 153
  - preview 153
  - subject 153
  - template 153
- message domain 54, 55, 183
- message domain attributes 55
- message domain name 54
- message domain properties 54
- message domains
  - for schedules 131
  - promoting 69
- message filters
  - for schedules 131
- message selector 183
- message template 186, 187
- message text 183
- message-based job 54
- message-based processing example 133
- message-based step 54
- message-driven step 183
  - defining properties 183
- metadata 45, 46
  - user-defined 27, 38, 39, 40, 41, 42, 43, 44
  - version properties 37
- metrics
  - for scoring performance 102
- MIME type filtering 57
- MIME types
  - promotion policies 57
- min function 233
- min/max
  - column attributes 209
- miscellaneous functions 243
- model evaluation 138, 139, 149, 185
  - notifications 149, 185
  - return codes 149, 185
- model execution date
  - filtering model management views 144
- model management 138, 139
- models
  - champion challenger 189
  - comparing 189
  - evaluating 189
  - IBM SPSS Modeler 189
  - monitoring 189
- modifying
  - permissions 28
  - server clusters 61
- modulo function 234
- month function 242
- monthly schedules 131
- mouse 9
- multiplication function 234

## N

- naming conventions 10
- naming factory 54, 55
- naming service 54, 55
- narrow search 18, 19, 20
- narrowing 18, 19, 20
- navigating through the system 203
- navigation 9, 203
  - enter key 9
  - mouse 9
- negation function 235
- Netezza 51, 204
- new
  - custom properties 38, 39
  - topics 42, 43
- new jobs 119
- nil values 229
- NOSPLASH 168
- NOSTATUSWIN 168
- not equal to function 245
- not function 249
- notification job step
  - adding to a job 185
  - body 186
  - connecting to other job steps 185
  - customizing 186
  - email attachments 185
  - from address 186
  - general information 186
  - iterative consumer 186
  - message template 186, 187
  - Notification tab 186
  - outcome-based 185
  - recipients 186
  - subject 186
  - use in model evaluation 185
- notifications 129
  - action 147
  - attachment 149, 150
  - body text 153
  - content 147
  - content notifications 148
  - defining 148
  - disabling 129
  - entering recipient email 150
  - errors 153
  - event property variables 153
  - folder events 148
  - folder notifications 148
  - folder options 149, 150
  - folder structure changes 148
  - for exported and imported objects 147
  - generating a list of recipients
    - dynamically 152, 153
  - HTML formatting 153
  - job 147
  - job step 148
  - label 147, 149, 150
  - message customization 149, 150, 153
  - message subject 149, 150
  - model evaluation 149
  - modifying 148
  - preview 149, 150, 153
  - previewing message 154
  - recipients 149, 150
  - removing 148, 150

- notifications (*continued*)
  - reverting to the default template 150
  - selecting recipients from Lotus Notes 151
  - selecting recipients from MS Outlook 151
  - selecting recipients from security subscribers 151
  - sender 149, 150
  - subject 153
  - template 149, 150, 153
- notifications delivery failure 155

## O

- object permissions
  - cascading 28
  - deleting 29
  - modifying 28
- object properties 25
  - editing 26
  - viewing 25, 26, 33, 37
- objects
  - content 13, 16, 17
  - deleting 17
  - files 16, 17
  - locking 22
  - permissions 16
  - search 17, 18, 19, 20, 21
  - unlocking 22
  - versions 37
- ODBC 118
- ODBC data sources 49
  - champion challenger 195
- ODS 169
- or function 249
- Outlook 151
- output 127
  - log files 138
- output file permissions 127
- output files
  - General job steps 174, 175
- output permissions 127
- overriding data sources
  - for data access plans 78, 79
- overview 203
  - content explorer 13
  - jobs 117, 118

## P

- parameter filters
  - for data access plans 79
- pass connector 124
- passwords
  - changing 15
  - modifying 15
  - new 15
- PDF settings
  - for the visually impaired 201
- permission inheritance 28
- permissions
  - applying 28
  - assigning 28
  - cascading 28
  - creating 28

- permissions (*continued*)
  - deleting 29
  - editing for label 34
  - for labels 34
  - for reporting results 159
  - for SAS results 168
  - for visualization results 164
  - inheritance 16
  - job step 127
  - modifying 28, 29
  - object 28, 29
  - output file 127
  - removing 29
  - specifying as iterative variables 128
  - updating in bulk 46
  - when importing 64
- PMML 95
- PMML files 175, 176
- policy 55, 56, 57, 69
- predictor effectiveness 140
- preferences
  - distribution channels 36
- prerequisites
  - jobs 118
- preview 149, 150
- preview button 231
- preview data
  - Data Provider Definition - Real Time 223
- preview of selection 231
- previewing notification message 154
- principals
  - adding for label 34
  - deleting for label 34
  - modifying label permissions 34
- privileges
  - access 38, 39, 42, 43
  - administrative 38, 39, 42, 43
- promotion 55, 56, 149
  - actions for roles 69
  - delayed 56, 57, 69
  - dependent object 57
  - immediate 56, 57
  - MIME types 57
  - policy 69
- promotion policies
  - adding 55
  - deleting 57
  - resource definitions 56
  - timing 56
- promotion policy
  - modifying 57
- properties
  - bulk updates 45, 46
  - content object 25
  - custom 27, 38, 39, 40, 41, 43
  - definition 25
  - deleting 20, 21
  - editing 20, 26
  - general 26
  - object 25, 26
  - overview 25
  - server 27, 33, 37, 38, 39, 40, 41, 42, 43, 44
  - topics 42, 43, 44
  - viewing 25, 26, 33, 37
- property 18, 19, 20

- property types
  - custom properties 27
- property variables 10

## R

- range of recurrence
  - for schedules 131
- reactivating
  - expired files 33
- real-time
  - data access plans 75, 79
- recipient
  - of notification 149, 150
- recipients 149, 150, 153
- recurrence pattern
  - for schedules 131
- reference filters
  - for data access plans 79
- references
  - external 63
- refine search 18, 19, 20, 21
- relationships
  - between steps 123, 124, 125
  - deleting 126
  - in data models 82
- remote process
  - execution servers 2, 5
- remote process servers
  - assigning to job steps 167, 173
  - server definition 59
- remotely-deployed scoring servers 6
- removing 150
  - categories 209
  - custom properties 41
  - locks 22
  - topics 44
- removing columns
  - Application View 212
  - IBM SPSS Collaboration and Deployment Services Enterprise View 207
- removing tables
  - Application View 212
  - IBM SPSS Collaboration and Deployment Services Enterprise View 206
- renaming
  - topics 43
- reorder search terms 20
- reordering 20
  - search terms 20
- reporting steps 157, 163
  - adding to jobs 157
  - burst 158
  - clean-up 160
  - data source 157
  - defining general properties 157
  - iterative 158
  - locale 159
  - metadata 159
  - naming 157
  - notifications 161
  - output file format 159
  - output file location 159
  - output file name 159
  - parameters 159

- reporting steps *(continued)*
  - prompts 159
  - rendering options 159
  - result permissions 159
  - results 159
  - single 158
  - type 158
  - variables 159
  - versions 157
- reports
  - external 197, 198
  - submitted 197, 198
- repository
  - connecting to 14, 15
  - content 13, 16, 17
  - deleting 17
  - disconnecting from 15
  - files 16, 17
  - locking objects 22
  - objects 22
  - search 17, 18, 19, 20, 21
  - unlocking objects 22
- resolving conflicts
  - globally 66
  - individually 66
  - individually 67, 68
- resource definitions
  - credential definitions 47, 48
  - data source definitions 47
  - data sources 48
  - exporting 62
  - importing 62
  - in promotion policies 56
  - locking 22
  - server definitions 47, 57, 58, 59
- resource definitions folder 13
- restrictions
  - expiration dates 32
  - exporting 64
  - importing 64
  - submitted jobs 197
- results 127, 138
  - analysis 138, 139, 140
  - champion challenger 138, 139
  - job 137, 138
  - job step 138
  - model evaluation 138, 139
  - predictor effectiveness 140
  - scenario 138, 139, 140
  - search 21
  - streams 138, 139
- resuming
  - scoring configurations 114
- return codes
  - model evaluation 149
  - notifications 149
- reverting to the default template 150
- rows
  - filtering 142, 143
- RSS feeds 36
- rtrim function 251
- running jobs 123, 129, 130, 136

## S

- SAS
  - execution server 2, 5, 59

- SAS *(continued)*
  - server definition 59
- SAS execution servers
  - assigning to job steps 167
- SAS steps 167
  - adding to jobs 167
  - controlling processing 167, 168, 169
  - creating graphs 168, 169
  - defining properties 167, 168
  - example 169
  - external files 169
  - HTML output 169
  - naming 167
  - naming results 168
  - result locations 168
  - result permissions 168
  - text output 169
  - versions 167
- save
  - Application View 215
  - Data Provider Definition 221
  - Data Provider Definition - Real Time 227
  - IBM SPSS Collaboration and Deployment Services Enterprise View 211
- saving jobs 126
- scenario
  - results 138, 139, 140
  - status 138, 139, 140
- schedule
  - job 142
- schedules 130
  - creating 130
  - credentials 130
  - deleting 133
  - dormant 133
  - editing 133
  - job labels 130
  - job variables 132
  - message-based 131
  - time-based 131
- scheduling jobs 129, 136
- schema browser 227
- scoring
  - A/B split testing 111, 112
  - configurations 98
  - models 95
  - performance 115
  - using PMML files 176
- scoring configurations 98, 114
  - aliases 105
  - association sets 107, 110, 111
  - auditing 102
  - batch scoring 104
  - cache sizes 104
  - configuration names 99
  - creating 98
  - data 99
  - data providers 99
  - deleting 114
  - editing 114
  - input 101
  - labels 99
  - logging 102
  - models 99
  - output 101

- scoring configurations *(continued)*
  - resuming 114
  - suspending 114
  - viewing 113
- scoring graph view 115
- scoring servers 6
- scoring view 113
  - deleting configurations 114
  - editing configurations 114
  - filtering 114
  - resuming configurations 114
  - suspending configurations 114
- screen elements 203
- screen readers 201
- scripting
  - champion challenger 189
  - IBM SPSS Modeler 189
- search 20, 21
  - accessing 18
  - advanced 17, 18, 19, 20, 21
  - AND 18, 19, 20
  - content objects 17, 18, 19, 20, 21
  - date range 18, 19, 20
  - deleting 21
  - dialog 18
  - expiration dates 33
  - files 17, 18, 19, 20, 21
  - group 20
  - OR 18, 19, 20
  - property 18, 19, 20
  - refining 18, 19, 20, 21
  - reordering 20
  - simple 17, 18, 19, 20, 21
  - time 18, 19, 20
  - ungroup 20
- search function 252
- search terms 18, 19
  - deleting 21
  - editing 20
  - group 20
  - properties 20, 21
  - reordering 20
  - stopwords 17, 21
  - ungroup 20
- searching
  - content objects 21
  - custom properties 40
  - files 21
  - for jobs 142, 143
  - submitted jobs 197
  - topics 44
- security
  - for labels 34
- security providers 128
- security subscribers 151
- select Application View
  - Data Provider Definition 218
  - Data Provider Definition - Real Time 222
- select data set to add
  - connection settings 226
- select tables to add
  - add data set 225
- selecting
  - versions 31
- selection
  - drag and drop 9

- selection functions 245
  - selection values
    - deleting 27, 40
    - editing 27, 40
    - modifying 27, 40
    - removing 27, 40
  - sender 149, 150
    - of notification 149, 150
  - sequential connector 124
  - server
    - change password 15
    - connection 33, 37
    - connections 33
    - custom properties 27, 38, 39, 40, 41
    - log in 14, 15
    - log off 14, 15
    - properties 33, 37
    - topics 27, 42, 43, 44
    - URL 33
    - version label 37
  - server clusters 60
    - adding servers 61
    - creating 60
    - modifying 61
    - name 61
    - promoting 69
    - removing servers 61
    - settings 61
    - weights 61
  - server connection 14
    - changing password 15
    - ending 15
    - existing 15
    - new 14
    - terminating 15
  - server connections
    - jobs 118
  - server definitions 47, 57
    - adding 58
    - content repository 58
    - location 58
    - modifying 59
    - names 58
    - remote process 59
    - SAS 59
    - types 58
  - server destination 58
  - server status 145
  - servers
    - promoting 69
    - remote process servers 167, 173
    - SAS 167
  - shortcuts
    - keyboard 199
  - simple search 17, 18, 19, 20, 21
  - size function 233
  - split testing
    - for scoring 111, 112
  - SSL 33
  - standard error 173
  - standard output 173
  - start date
    - for schedules 131
  - start time
    - for schedules 131
  - starting
    - application 9
  - starting (*continued*)
    - client 9
    - system 9
  - startsWith function 252
  - status
    - filtering 142, 143
    - job 137
    - job step 138
    - server 145
  - stderr 173
  - stdout 173
  - steps
    - adding 122
    - executing 123
    - job 117, 118, 119, 120, 122, 123
    - relationships 123, 124, 125, 126
  - str function 239
  - streams
    - results 138, 139
    - status 138, 139
  - string function 252
  - string functions 250
  - subject 149, 150
  - submitted jobs
    - expiration dates 198
    - restrictions 197
    - searching 197
  - subscribing to files 154
  - subscription delivery failure 155
  - subscriptions
    - email address for 155
    - files 154
    - for exported and imported objects 147
    - jobs 154
    - managing 155
    - modifying 154
    - removing 155
    - subscribing to files 154
    - unsubscribing 154
    - user preferences 155
  - subtraction function 235
  - sum function 234
  - suspending
    - scoring configurations 114
  - syndication feeds 36
  - system
    - content explorer 13
    - exiting 11
    - help 9
    - jobs 117
    - launching 9
    - naming files 10
    - navigating 9
    - starting 9, 14
  - T**
    - table
      - search results 21
      - select columns 205
    - table properties 53
    - tables 53
      - Application View 212
      - champion challenger 138, 139
      - filtering 142, 143
  - tables (*continued*)
    - IBM SPSS Collaboration and Deployment Services Enterprise View 205
    - job history 137, 143
    - job schedule 142
    - job step history 137, 138
    - keyboard navigation 200
    - model evaluation 138, 139
    - predictor effectiveness 140
    - server status 145
  - tables keys 53
  - template 149, 150
  - Teradata 51, 204
  - terms
    - deleting 21
    - editing 20
    - group 20
    - reordering 20
    - search 18, 19, 20
    - search exclusions 17, 21
    - ungroup 20
  - text output
    - for SAS steps 169
  - time 18, 19, 20
  - timeout 183
  - timestamp function 240
  - toLower function 253
  - topics
    - accessing 43
    - copying 43
    - creating 27, 42, 43
    - deleting 44
    - editing 27
    - existing 27
    - label 43
    - modifying 27
    - new 43
    - property type 43
    - renaming 43
    - searching 44
    - selection values 27, 43
  - toUpper function 253
  - trend
    - filtering model management views 144
  - type field 229
- U**
- UI 203
  - UNC file references 118
  - ungroup search terms 20
  - unlocking 22
  - unlocking objects 22
  - unmap columns 219, 222
  - use versions permission
    - for labels 34
  - user interface 203
  - user preferences 4
    - distribution channels 36
  - email 155
  - users 128
    - deleting 29
    - existing 28
    - new 28
    - permissions 127

## V

- validate
  - Data Provider Definition 220
- variables
  - adding to jobs 121
  - editing 122
  - in jobs 121, 129
  - removing from jobs 122
- version
  - description 37
  - expiration date 37
  - keywords 37
  - metadata 37
  - properties 37
  - values 37
- version label
  - filtering 142, 143
- version properties
  - updating in bulk 45
- versions
  - deleting 31
  - invalid 67, 68
  - labeling 37
  - labels 29, 30, 31
  - promoting 69
  - selecting 31
- viewing
  - expired files 33
  - job history 137, 143
  - job output 138
  - job schedule 142
  - job step history 137, 138
  - logs 138
  - properties 25, 26, 33, 37
  - search results 21
  - server status 145
- views
  - champion challenger 138, 139
  - model evaluation 138, 139
  - model management 138, 139
  - predictor effectiveness 140
  - server status 145
- visualization reports
  - adding to jobs 163
  - clean-up 165
  - data source 163
  - defining general properties 163
  - dimensions 164
  - iterative 164
  - metadata 164
  - naming 163
  - notifications 166
  - output file location 164
  - parameters 164
  - prompts 164
  - result permissions 164
  - results 164
  - single 164
  - type 164
  - variables 164
  - versions 163
- visualization reports steps
  - output file format 164

## W

- weekofyear function 243
- weely schedules 131
- weights
  - in server clusters 61
- when function 244
  - handling nil values 229
- working directory
  - for General job steps 173, 174, 175
- working with data sets
  - Data Provider Definition - Real Time 224

## X

- XOM
  - See* execution object models

## Y

- year function 243









Printed in USA