

SAP BusinessObjects Business Intelligence platform
Document Version: 4.1 Support Package 3 - 2014-05-07

Information Design Tool User Guide



Table of Contents

1	What's new in the Information Design Tool User Guide.	14
2	Getting started with the information design tool.	18
2.1	About the information design tool.	18
2.2	About resources in the information design tool.	19
2.3	Starting the information design tool.	21
2.4	About the information design tool interface.	22
2.4.1	Resetting the user interface display.	24
2.5	Setting preferences in the information design tool.	24
2.5.1	Setting preferences for the Business Layer Editor.	25
2.5.2	Setting preferences for check integrity	26
2.5.3	Setting connection display preferences for the Data Foundation Editor.	27
2.5.4	Setting display preferences for the data foundation view.	28
2.5.5	Setting table and join detection options.	29
2.5.6	Setting default join states for contexts.	31
2.5.7	Setting performance-related options for the data foundation view.	31
2.5.8	Setting languages used by the information design tool.	32
2.5.9	Setting a link for online tutorials.	33
2.5.10	Setting middleware for secured relational connections.	33
2.5.11	Setting preferences for showing values.	34
2.6	How to get help using the information design tool.	35
2.7	Wizards in the information design tool.	36
3	Creating Universes.	37
3.1	How to create a universe with the information design tool	37
3.2	Using SAP NetWeaver BW data sources.	40
3.2.1	How InfoProvider objects are mapped in a universe.	41
3.2.2	Refreshing universes based on SAP Netweaver BW.	44
3.3	Using SAP HANA data sources.	45
3.3.1	Creating relational resources on SAP HANA information models.	47
3.4	Using SAP ERP data sources.	48
3.5	Using Microsoft Analysis Services (MSAS) data sources.	49
3.6	Using Essbase data sources.	51
3.7	Using SAS data sources.	53
3.8	Multilingual universes.	53
3.8.1	Translating universe metadata.	55
3.9	About the New Universe wizard.	56
3.9.1	Select or create a project in the New Universe wizard.	56
3.9.2	Select the data source type in the New Universe wizard.	57

3.9.3	Select or create a relational connection in the New Universe wizard.	57
3.9.4	Select or create an OLAP connection in the New Universe wizard.	57
3.9.5	Select or create a data foundation in the New Universe wizard.	58
4	Converting .unv universes.	59
4.1	About .unv and .unx universes.	59
4.2	About converting .unv universes	59
4.3	Features supported when converting .unv universes.	61
4.4	Tips for resolving check integrity errors after converting .unv universes.	65
4.5	Converting a .unv universe in a repository.	67
4.6	Converting a locally-stored .unv universe.	68
5	Retrieving published universes.	70
5.1	Retrieving a published universe from the local file system.	70
5.2	Retrieving a published universe from a repository.	70
6	Migrating universes to SAP HANA.	72
6.1	About Universe Landscape Migration.	72
6.1.1	Universe Landscape Migration requirements and limitations.	73
6.2	Migrating a universe to SAP HANA: Pre-Migration.	74
6.3	Migrating a universe to SAP HANA: Migration.	75
6.4	Migrating a universe to SAP HANA: Post-Migration.	76
7	Working with projects.	78
7.1	About local projects and resources.	78
7.1.1	Creating a local project.	79
7.1.2	About resource names.	79
7.1.3	Finding universe resources in the local file system.	80
7.1.4	Opening a local project.	80
7.1.5	Deleting a local project.	81
7.1.6	Backing-up and recovering universe resources in local projects.	81
7.1.7	Searching for and filtering resources in the Local Projects View.	82
7.2	About shared projects.	82
7.2.1	Creating a shared project from a local project.	83
7.2.2	Working in a shared project.	84
7.2.3	Renaming a shared project.	84
7.2.4	Deleting a shared project.	85
7.3	About project synchronization.	86
7.3.1	Opening the Project Synchronization View.	88
7.3.2	Synchronizing a project.	89
7.3.3	Locking a resource.	90
7.3.4	Unlocking a resource.	91
7.3.5	Merging changes to shared resources.	92

7.4	Saving resources as reports.	93
8	Working with repository resources.	94
8.1	About managing repository resources.	94
8.2	About session management.	95
8.2.1	Opening a session.	96
8.2.2	Closing a session.	97
8.3	Running a query on a universe published in a repository.	98
9	Working with connections.	99
9.1	About connections.	99
9.1.1	About local connections.	100
9.1.2	About secured connections.	100
9.1.3	About connection shortcuts.	101
9.2	About the Connection Editor.	102
9.3	Creating a relational connection.	102
9.3.1	Name a connection.	103
9.3.2	Select a middleware driver.	104
9.3.3	Set the connection parameters.	104
9.4	Creating an OLAP connection.	123
9.4.1	Select an OLAP middleware driver.	124
9.4.2	Set login parameters for OLAP data sources.	124
9.4.3	Select an OLAP cube	129
9.5	Creating a connection shortcut.	129
9.6	Editing local and secured connections.	130
9.7	Editing connection shortcuts.	131
9.8	Showing values in a relational connection.	131
9.9	Showing values in an OLAP connection.	132
10	Working with data foundations.	133
10.1	About data foundations.	133
10.1.1	About data foundation types.	133
10.1.2	About single-source data foundations.	133
10.1.3	About multisource-enabled data foundations.	134
10.2	About the Data Foundation Editor.	136
10.3	How to build a data foundation.	138
10.4	About connections in the data foundation.	141
10.4.1	Adding connections to a data foundation.	143
10.4.2	Changing a connection in a data foundation.	143
10.4.3	Selecting delimitation overrides to keep.	144
10.4.4	Searching for tables in the Connection pane.	145
10.4.5	Filtering the tables in the connection by table type.	147

10.4.6	Filtering tables in an SAP HANA connection by information model.	148
10.5	About tables in the data foundation.	148
10.5.1	Inserting tables into the data foundation.	151
10.5.2	Editing table properties.	152
10.5.3	Setting case of table names.	154
10.5.4	Hiding and unhiding table columns.	154
10.5.5	Changing column data types.	155
10.5.6	Changing qualifiers and owners.	155
10.5.7	Changing table and column delimitation.	156
10.6	About table keys.	157
10.6.1	Setting and detecting table keys.	158
10.7	About table row counts.	159
10.8	About joins.	160
10.8.1	Inserting and editing a join.	161
10.8.2	Detecting joins.	163
10.8.3	Inserting a column filter.	164
10.9	About cardinality.	165
10.9.1	Detecting and setting cardinalities.	165
10.10	Inserting a calculated column.	166
10.11	Inserting a time column.	168
10.12	About derived tables.	168
10.12.1	Inserting a derived table based on a data foundation table.	169
10.12.2	Merging tables.	170
10.12.3	Inserting and editing a derived table.	171
10.13	About alias tables.	172
10.13.1	Inserting alias tables.	172
10.13.2	Detecting alias tables.	173
10.13.3	Highlighting aliases.	174
10.13.4	Highlighting the original table of an alias.	174
10.14	About contexts.	175
10.14.1	Detecting contexts.	176
10.14.2	Inserting and editing contexts.	177
10.15	Resolving loops.	177
10.16	About input columns in the data foundation.	178
10.16.1	Editing input columns.	179
10.17	About parameters and lists of values in the data foundation.	180
10.18	About data foundation properties.	180
10.18.1	Editing SQL options in the data foundation.	181
10.18.2	Setting SQL generation parameters in the data foundation.	182
10.18.3	Showing a data foundation summary.	183
10.18.4	Editing the data foundation description and comments.	183

10.18.5	Showing SAP HANA variable information.	184
10.19	Showing table values.	184
10.19.1	Showing values in a data source.	185
10.20	Showing column values.	186
10.21	Profiling column values.	187
10.22	Showing local dependencies in the data foundation.	187
10.23	About refreshing a data foundation.	188
10.23.1	Synchronizing tables.	189
10.24	Inserting a custom data foundation view.	190
10.25	Searching for tables and columns in the data foundation.	191
10.26	Inserting a comment into the data foundation view.	192
10.27	Centering the view on a selection.	192
10.28	Changing the display of objects in the data foundation.	193
10.28.1	Auto-arranging tables in the data foundation view.	193
10.28.2	Changing table display.	194
10.28.3	Grouping tables using families.	194
11	Working with the federation layer.	196
11.1	About the federation layer.	196
11.2	Building the federation data flow.	196
11.3	About federated tables.	198
11.3.1	Adding a federated table manually.	199
11.3.2	Adding a federated table from a data source.	200
11.3.3	Editing a federated table.	200
11.4	About input tables and joins.	201
11.4.1	Adding input tables to a mapping.	202
11.4.2	Joining input tables.	203
11.4.3	Configuring meanings of input table joins using core tables.	204
11.5	About mappings in the federation layer.	205
11.5.1	Mapping columns from the input table to columns of the federated table.	206
11.5.2	Editing a mapping formula.	206
11.5.3	About the SQL Expression Editor.	207
11.5.4	Adding a mapping.	207
11.5.5	Activating and deactivating mappings.	208
11.6	About distinct rows on input tables.	209
11.6.1	Activating and deactivating distinct rows.	209
11.7	About pre-filters and post-filters.	209
11.7.1	Adding and editing pre-filters.	210
11.7.2	Editing post-filters.	210
11.8	Showing values in a federated table.	211
11.9	Checking integrity of the federation layer.	212
11.10	Inserting a federated table into the data foundation.	212

11.11	Refreshing the structure of the federation layer.	213
12	Working with business layers.	215
12.1	About business layers.	215
12.2	About business layer objects.	215
12.3	How to build a relational business layer.	218
12.3.1	Specifying the type of data source for a business layer.	220
12.3.2	Naming a business layer.	220
12.3.3	Selecting a data foundation for a business layer.	221
12.4	How to build an OLAP business layer.	222
12.4.1	Selecting an OLAP connection and cube for a business layer.	224
12.4.2	Selecting an Essbase Accounts dimension.	224
12.4.3	Selecting objects from an OLAP cube for a business layer.	225
12.5	About the Business Layer Editor.	225
12.5.1	Changing display options of the business layer tree view.	226
12.6	About business layer properties.	227
12.6.1	OLAP data source properties.	228
12.6.2	About query stripping.	229
12.6.3	Editing the business layer name, description, and comments.	230
12.6.4	Editing query limits and options in the business layer.	231
12.6.5	Changing the data source of a business layer.	231
12.6.6	Setting SQL generation parameters in the business layer.	232
12.6.7	Displaying a business layer summary.	233
12.7	About index awareness.	234
12.8	About analytic functions.	234
12.8.1	Analytic functions: syntax and examples.	235
12.8.2	Analytic functions: rules, restrictions, and best practices.	238
12.8.3	Using analytical functions in a business layer object definition.	239
12.8.4	Using analytical functions in a derived table definition.	240
12.9	About aggregate awareness.	241
12.9.1	Setting up aggregate awareness.	241
12.9.2	Setting aggregate navigation.	242
12.10	Working with business layer objects.	243
12.10.1	Inserting a folder.	243
12.10.2	Inserting and editing dimensions.	244
12.10.3	Inserting dimensions directly from the data foundation.	246
12.10.4	Defining keys for dimensions and dimension attributes.	247
12.10.5	Turning an attribute or measure into a dimension.	248
12.10.6	Inserting and editing measures.	249
12.10.7	Turning a dimension or attribute into a measure.	251
12.10.8	Inserting and editing attributes.	253
12.10.9	Turning a dimension or measure into an attribute.	254

12.10.10	Inserting and editing filters.	255
12.10.11	Inserting and editing analysis dimensions.	257
12.10.12	Inserting and editing hierarchies.	258
12.10.13	Inserting and editing hierarchy levels.	260
12.10.14	Inserting and editing named sets.	261
12.10.15	Inserting and editing calculated members.	262
12.10.16	Defining the SQL expression for an object.	264
12.10.17	Defining the MDX expression for an object.	264
12.10.18	Associating extra tables.	265
12.10.19	Changing the state of an object: Active, Hidden, or Deprecated.	266
12.10.20	Setting object access levels.	267
12.10.21	Setting where objects can be used.	267
12.10.22	Setting options for the default list of values.	268
12.10.23	Creating and editing display formats for business layer objects.	269
12.10.24	About source information for business layer objects	276
12.10.25	Inserting and editing custom properties.	276
12.10.26	Showing associated objects.	277
12.10.27	Showing business layer object values.	277
12.10.28	Searching for business layer objects.	278
12.11	About business layer views.	279
12.11.1	Creating and editing a business layer view.	279
12.11.2	Filtering by business layer view.	280
12.12	About parameters	280
12.12.1	Inserting and editing a parameter.	281
12.12.2	Creating an index-aware prompt.	282
12.13	About lists of values	283
12.13.1	Inserting or editing a list of values	284
12.13.2	List of values column properties.	286
12.13.3	Associating a list of values with a business object.	287
12.13.4	Associating a list of values with a prompt defined in the business layer.	287
12.14	About navigation paths for objects.	288
12.14.1	Inserting a navigation path object into a business layer.	289
12.15	About queries in a business layer.	290
12.15.1	Inserting and editing a query in the business layer.	290
12.16	Reordering objects in the Business Layer Editor.	291
12.17	About refreshing business layers.	291
12.17.1	Refreshing an OLAP business layer	293
12.17.2	Inserting candidate objects.	293
12.18	About computing statistics for optimized query execution	294
12.18.1	Computing statistics for a multisource-enabled universe.	295

13	Using the Query Panel.	297
13.1	How to build a query.	297
13.2	About the Member Selector.	298
13.2.1	About selecting hierarchy members.	299
13.2.2	Opening the Member Selector in the Query Panel	300
13.2.3	Selecting hierarchy members.	300
13.2.4	Selecting members by hierarchy relationship	301
13.2.5	Selecting hierarchy members by level.	302
13.2.6	Selecting named sets.	303
13.2.7	Selecting calculated members.	303
13.2.8	Searching for hierarchy members.	304
13.2.9	Excluding hierarchy members.	305
13.2.10	Defining a prompt to select members.	305
13.2.11	Showing selected members in the Member Selector.	306
13.2.12	Sorting hierarchy members.	307
13.2.13	Setting display options.	307
13.2.14	Showing estimated child count.	307
13.3	Filtering data in the Query Panel.	308
13.3.1	How to build a business filter.	308
13.3.2	Filtering data using prompts.	310
13.4	Setting query properties.	313
13.5	Viewing and editing the query script.	314
13.6	Profiling column values in the query panel.	315
14	Checking integrity.	316
14.1	Running check integrity.	316
14.2	Reviewing check integrity problems.	317
15	Showing dependencies between resources.	318
15.1	About resource dependencies.	318
15.2	Showing local dependencies.	320
15.3	Showing repository dependencies.	321
16	Publishing resources.	323
16.1	About publishing resources.	323
16.2	Publishing a universe.	324
16.2.1	Selecting a repository folder.	325
16.2.2	Selecting a local folder.	325
16.3	Publishing a local connection to the repository.	326
16.4	Publishing a local universe to the repository.	326
17	Managing security.	328
17.1	About universe security.	328

17.2	About securing resources in the information design tool.	329
17.3	CMC rights for information design tool users.	330
17.4	About the Security Editor.	332
17.5	How to secure a universe using security profiles.	333
17.6	Opening the Security Editor.	335
17.7	Inserting and editing a Data Security Profile.	336
17.7.1	Data Security Profile settings.	337
17.7.2	Data Security Profile Connections settings.	338
17.7.3	Data Security Profile Controls settings.	338
17.7.4	Data Security Profile SQL settings.	339
17.7.5	Data Security Profile Rows settings.	340
17.7.6	Data Security Profile Tables setting.	341
17.8	Changing Security Profile priority.	341
17.9	Inserting and editing a Business Security Profile.	342
17.9.1	Business Security Profile settings.	343
17.9.2	Business Security Profile Connections settings.	344
17.9.3	Business Security Profile Create Query settings.	345
17.9.4	Business Security Profile Display Data settings.	346
17.9.5	Business Security Profile Filters settings.	347
17.10	Security profile aggregation.	348
17.10.1	Aggregation of Connections settings.	349
17.10.2	Aggregation of Controls settings.	349
17.10.3	Aggregation of SQL settings.	350
17.10.4	Aggregation of Rows settings.	351
17.10.5	Aggregation of Tables settings.	352
17.10.6	Aggregation of Create Query settings.	352
17.10.7	Aggregation of Display Data settings.	353
17.10.8	Aggregation of Filters settings.	354
17.11	Changing security profile aggregation options.	355
17.12	Assigning Security Profiles to users.	356
17.13	Displaying profiles assigned to a user and previewing net profiles.	356
18	SQL and MDX reference.	358
18.1	About the SQL/MDX Expression Editor.	358
18.2	SAP BusinessObjects SQL function reference for multisource-enabled universes	360
18.2.1	Aggregation functions.	360
18.2.2	ASCII Code (ascii).	364
18.2.3	Absolute (abs).	365
18.2.4	Angle Tangent 2 (atan2).	365
18.2.5	Arc Tangent (atan).	366
18.2.6	Arc Cosine (acos).	367

18.2.7	Arc Sine (asin).....	367
18.2.8	Case.....	368
18.2.9	Cast.....	370
18.2.10	Catalog.....	371
18.2.11	Ceil (ceiling).....	371
18.2.12	Character (char).....	372
18.2.13	Charindex (pos) (locate).....	372
18.2.14	Concat.....	373
18.2.15	Contains Only Digits.....	374
18.2.16	Convert.....	375
18.2.17	Cosine (cos).....	376
18.2.18	Cotangent (cot).....	376
18.2.19	Current Date (curDate).....	377
18.2.20	Current Time (curTime).....	377
18.2.21	Database.....	377
18.2.22	Day Name.....	378
18.2.23	Day Of Month.....	379
18.2.24	Day Of Week.....	379
18.2.25	Day Of Year.....	380
18.2.26	Decrement Days.....	380
18.2.27	Degrees.....	381
18.2.28	Exp.....	381
18.2.29	Floor.....	382
18.2.30	Hexa To Int.....	382
18.2.31	Hour.....	383
18.2.32	If Else.....	384
18.2.33	If Null (nvl).....	384
18.2.34	Increment Days.....	385
18.2.35	Int To Hexa.....	385
18.2.36	Is Like.....	386
18.2.37	LPad.....	387
18.2.38	Left.....	388
18.2.39	Left Remove (ltrim).....	389
18.2.40	Length.....	389
18.2.41	Log.....	390
18.2.42	Log10.....	391
18.2.43	Lowercase (lcase).....	391
18.2.44	Minute.....	392
18.2.45	Mod.....	392
18.2.46	Month Name.....	393

18.2.47	Now.....	394
18.2.48	Number of the Month (month).....	394
18.2.49	Number of the Week (week).....	395
18.2.50	Permute.....	395
18.2.51	Pi.....	397
18.2.52	Power.....	397
18.2.53	Quarter.....	398
18.2.54	Radians.....	398
18.2.55	Random (rand).....	399
18.2.56	Replace.....	399
18.2.57	Replace String Exp.....	400
18.2.58	Replicate (repeat).....	401
18.2.59	Rightpart (right).....	402
18.2.60	Round.....	402
18.2.61	Rpad.....	403
18.2.62	Rpos.....	404
18.2.63	Rtrim.....	405
18.2.64	Schema.....	405
18.2.65	Second.....	406
18.2.66	Sign.....	406
18.2.67	Sine (sin).....	407
18.2.68	Space.....	407
18.2.69	Sqrt.....	408
18.2.70	Stuff (insert).....	408
18.2.71	Substring.....	409
18.2.72	Tangent (tan).....	410
18.2.73	Timestamp Add.....	411
18.2.74	Timestamp Diff.....	412
18.2.75	To Boolean.....	413
18.2.76	To Date.....	414
18.2.77	To Decimal.....	415
18.2.78	To Double.....	415
18.2.79	To Integer.....	416
18.2.80	To Null.....	416
18.2.81	To String.....	417
18.2.82	To Time.....	418
18.2.83	To Timestamp.....	419
18.2.84	Trim.....	420
18.2.85	Trunc.....	420

18.2.86	Uppercase (ucase).....	421
18.2.87	User.....	422
18.2.88	Year.....	422
18.3	About @Functions.....	423
18.3.1	About @Aggregate_Aware.....	423
18.3.2	About @DerivedTable.....	424
18.3.3	About @Execute.....	424
18.3.4	About @Prompt.....	426
18.3.5	About @Select.....	430
18.3.6	About @Variable.....	430
18.3.7	About @Where.....	432
18.4	About SQL Generation Parameters.....	433
18.4.1	SQL generation parameters reference.....	433
18.4.2	SQL generation parameters set in the extended PRM.....	445

1 What's new in the Information Design Tool User Guide

Links to information about the new features and documentation changes for the information design tool for each version of SAP BusinessObjects BI platform.

SAP BusinessObjects BI platform 4.1 Support Package 3 - March 2014

What's new	Link to more information
Enhanced query stripping method for relational universes.	About query stripping [page 229]
DELIMITER parameter for the @Variable function.	About @Variable [page 430]
Clarified table types allowed in table replacements in Data Security Profiles.	Data Security Profile Tables setting [page 341]

SAP BusinessObjects BI platform 4.1 Support Package 2 - November 2013

What's new	Link to more information
An application preference that lets you change the default state for joins in existing contexts when joins are added to the data foundation. You can also choose to use the new default behavior when adding contexts.	Setting default join states for contexts [page 31]
For connections to Oracle data sources using JDBC middleware, a configuration parameter Query Timeout that lets you modify the number of seconds before a query times out.	Configuration parameters for relational connections [page 116]
Added information on how to use analytic functions in the information design tool, including examples.	About analytic functions [page 234]

SAP BusinessObjects BI platform 4.1 Support Package 1 - August 2013

What's new	Link to more information
Added information on how to create an index-aware prompt.	Creating an index-aware prompt [page 282]

What's new	Link to more information
Removed information about two options for lists of values: Allow users to edit list of values and Automatic refresh before use . These options do not apply to universes created with the information design tool.	
Updated information about SAP HANA variables, showing values, and business layer creation on SAP HANA relational connections. Certain features apply only to single-source data foundations.	Using SAP HANA data sources [page 45]

SAP BusinessObjects BI platform 4.1 - May 2013

What's new	Link to more information
A federation layer in multisource-enabled data foundations. The federation layer lets you create federated tables that can include data from any of the data source connections defined in the data foundation.	About the federation layer [page 196]
Enhancements to relational universes on SAP HANA, including support for SAP HANA variables, and a wizard to automatically create a data foundation and business layer based on selected SAP HANA information models.	Using SAP HANA data sources [page 45]
The Universe Landscape Migration plug-in, which lets you to migrate existing relational universes to SAP HANA.	About Universe Landscape Migration [page 72]
Help for new information design tool users, including a welcome page, a cheatsheet to help create an OLAP universe, and a New Universe wizard.	How to get help using the information design tool [page 35] About the New Universe wizard [page 56]
Columns can be hidden in standard tables in the data foundation (supports features in new data sources).	About tables in the data foundation [page 148]
You can override the default delimitation for column names in the data foundation.	About tables in the data foundation [page 148]
You can enter empty values for input columns in the data foundation with character data type.	Editing input columns [page 179]
You can filter by table type or information model (SAP HANA connections) when inserting tables into the data foundation.	Inserting tables into the data foundation [page 151]
You can filter by column name when editing joins.	Inserting and editing a join [page 161]

What's new	Link to more information
Wildcard searches in the data foundation connection panel are now supported.	Searching for tables in the Connection pane [page 145]
You can filter by information model in the data foundation connection panel.	Filtering tables in an SAP HANA connection by information model [page 148]
Dimension attributes in the business layer can be index aware.	Defining keys for dimensions and dimension attributes [page 247]
Added information on index awareness.	About index awareness [page 234]
You can define attributes for measures in OLAP business layers.	Inserting and editing attributes [page 253]
Query stripping is available for relational and OLAP universes.	About query stripping [page 229]
You can share custom display formats for business layer objects between business layers.	Creating and editing display formats for business layer objects [page 269]
You can create a custom ordering of objects such as parameters and lists of values.	Reordering objects in the Business Layer Editor [page 291]
Command to turn measures or dimensions into attributes in the business layer.	Turning a dimension or measure into an attribute [page 254]
Command to turn dimensions or attributes into measures in the business layer.	Turning a dimension or attribute into a measure [page 251]
You can drag and drop table columns to create dimension and attribute keys in the business layer.	Defining keys for dimensions and dimension attributes [page 247]
Connections for OData data sources.	Parameters for OData connections [page 113]
Connections for XML and Web Services data sources.	Parameters for XML and Web Services connections [page 114]
Direct access connections for SAP HANA for SAP Crystal Reports for Enterprise.	Using SAP HANA data sources [page 45]
Additional SAP HANA connection parameters for relational and direct access.	Login parameters for relational connections [page 105] Login parameters for OLAP connections [page 125]
Additional relational connection parameters for SAP NetWeaver BW.	Login parameters for SAP NetWeaver BW and ERP connections [page 107]
Security profile setting for OLAP replacement connections.	Business Security Profile Connections settings [page 344]

What's new	Link to more information
Enhanced join support and naming of business layer objects for relational universes on SAP ERP.	Using SAP ERP data sources [page 48]

2 Getting started with the information design tool

2.1 About the information design tool

The information design tool is an SAP BusinessObjects metadata design environment that enables a designer to extract, define, and manipulate metadata from relational and OLAP sources to create and deploy SAP BusinessObjects universes.

A universe is an organized collection of metadata objects that enable business users to analyze and report on corporate data in a non-technical language. These objects include dimensions, measures, hierarchies, attributes, pre-defined calculations, functions, and queries. The metadata object layer, called the business layer, is built on a relational database schema or an OLAP cube, so the objects map directly to the database structures via SQL or MDX expressions. A universe includes connections identifying the data sources so queries can be run on the data.

The role of the universe is to provide the business user with semantically understandable business objects. The user is free to analyze data and create reports using relevant business language regardless of the underlying data sources and structures.

Universes created using the information design tool can be used by the following SAP data analysis and reporting applications starting with version BI 4:

- SAP BusinessObjects Web Intelligence
- SAP Crystal Reports for Enterprise
- SAP BusinessObjects Explorer
- SAP BusinessObjects Dashboard Design
- SAP Lumira
- SAP Predictive Analysis
- SAP Design Studio

i Note

Check the documentation for an application for any restrictions on accessing universes.

To enable the designer to create universes, the information design tool provides the resources necessary to do the following:

- Create connections to data sources.
- Extract a complete OLAP cube schema.
- Extract tables and joins to build a relational schema called a data foundation.
- Create metadata objects from the cube or the data foundation. These objects are contained and organized in a business layer. The SQL and MDX expressions within objects can be validated and queries run against the target databases to test the business layer.
- Share resources to allow multiple designers to work on the same resources concurrently.
- Publish a universe, which compiles the business layer, the data foundation, and the connections into a single universe file (.unx):
 - Publish a universe to a repository to be implemented in deployments of SAP BusinessObjects data analysis and reporting applications.

- Publish a universe locally, to be implemented by client applications in standalone mode (for example Web Intelligence Rich Client).
- Create security profiles to define user access to universe data and metadata.

Who uses the information design tool?

The universe designer may be a database administrator, an applications manager or developer, a project manager, or a report creator who has acquired enough technical skills to create universes for other users. A security administrator also uses the information design tool to define universe security profiles.

There can be more than one universe designer in a company. The number of universe designers depends on the company's data requirements. For example, one universe designer could be appointed for each application, project, department or functional area.

Related Information



[About resources in the information design tool](#) [page 19]








[Starting the information design tool](#) [page 21]




[How to create a universe with the information design tool](#) [page 37]

2.2 About resources in the information design tool

The information design tool provides the following design resources to extract metadata and build universes.

Resource	Description
 Project	<p>A project is a named local workspace that contains the resources used to build one or more universes.</p> <p>A project can be shared so that multiple designers can work on the same resources.</p> <p>A project can contain any number of independent resources, for example data foundations, business layers, and connections. All resources contained within a project can be used interchangeably, for example a connection can be used by several data foundations within the same project.</p> <p>Projects and their resources are displayed in the Local Projects View. To open a resource in the editor, double-click the resource in the Local Projects View.</p>
Connection  OLAP	<p>A connection is a named set of parameters that define how a universe can access a relational or OLAP data source. A universe is always associated with at least one connection. A connection is an independent resource and can be used by several</p>

Resource	Description
 Relational	<p>universes. You can build a multisource-enabled universe that references one or more relational connections.</p> <p>Connections can be local (stored in a local file) or secured (an object in a shared repository that is referenced by a connection shortcut).</p> <p>Local connections are stored in the local project as .cnx files.</p>
 Connection shortcut	<p>A connection shortcut is an object in the local project that references a secured connection in a repository. You use a connection shortcut to refer to secure connections when creating data foundations and business layers based on secure connections.</p> <p>Connection shortcuts are stored in the local project as .cns files.</p>
 Data foundation	<p>A data foundation is a schema that defines the relevant tables and joins from one or more relational databases. You enhance the data foundation by adding federated tables, derived tables, alias tables, calculated columns, additional joins, contexts, prompts, lists of values, and other SQL definitions. The data foundation becomes the basis of one or more business layers.</p> <p>Data foundations are stored in the local project as .dfx files.</p>
 Business layer	<p>A business layer is a collection of metadata objects that provides an abstraction of relational database entities or OLAP cubes, understandable by a business user. Objects map via SQL expressions to an underlying data foundation, or via MDX expressions to an underlying OLAP cube. These objects include dimensions, hierarchies, measures, attributes, and predefined conditions.</p> <p>You can add dimensions, hierarchies, measures, attributes, and other objects as the universe design requires. You can validate the SQL or the MDX at any time. You can create queries, lists of values, parameters (also called prompts), and navigation path objects.</p> <p>The business layer is the universe under construction, and when the business layer is complete, it is compiled with the connections or connection shortcuts and data foundation, published, and deployed as a universe.</p> <p>Business layers are stored in the local project as .blx files.</p>
 Query	<p>A query is a set of objects that define a request to the database for data. A query can be defined and saved in the business layer as a metadata object to be used to test objects in the business layer.</p>
 Parameter  List of values	<p>A parameter is a variable in the universe that requires a value at query time. Parameters are often defined to prompt the user to supply a value, and in this case are referred to as prompts.</p> <p>A list of values is a collection of data values that can be associated with an object in the universe, allowing the user to choose values for a prompt.</p> <p>Parameters and lists of values can be defined in the data foundation. They are inherited by all business layers based on that data foundation.</p>

Resource	Description
	Parameters and lists of values can also be defined in the business layer.
 Universe	<p>A universe is a compiled file that includes all resources used in the definition of the metadata objects built in the design of the business layer.</p> <p>The universe is used by SAP BusinessObjects data analysis and reporting applications, where the business layer objects are visible for analysis and reporting.</p> <p>Universes are stored either locally or in a repository as .unx files.</p>
Security Profiles  Data  Business	<p>A security profile is a group of security settings that controls the data and metadata that are displayed to users and modifies the parameters defined in the data foundation and/or business layer. Security profiles are defined on published universes and stored in the repository.</p>

Related Information

[How to create a universe with the information design tool](#) [page 37]

[About local projects and resources](#) [page 78]

[About connections](#) [page 99]

[About data foundations](#) [page 133]

[About business layers](#) [page 215]

[About queries in a business layer](#) [page 290]

[About parameters](#) [page 280]

[About lists of values](#) [page 283]

[About universe security](#) [page 328]

2.3 Starting the information design tool

Context

The information design tool is installed with the SAP BusinessObjects Business Intelligence platform Client Tools. For more information on installing the BI platform Client Tools, see the *SAP BusinessObjects Business Intelligence Suite 4.0 Master Guide*, or the *SAP Crystal Server 2011 Getting Started Guide*.

Once the client tools are installed on your machine, for example in a Windows installation of the BI platform, you can start the information design tool with the command: **Start > All programs > SAP Business Intelligence > SAP BusinessObjects BI platform 4 Client Tools > Information Design Tool**.

No authentication is required to use the information design tool in offline mode (not connected to a repository). You can begin creating and editing local resources.

Note

When you double-click a resource file in the local file system (for example a .blx, .dfx, or .cnx file), the information design tool opens without opening the specific resource editor. You must open the editor from the Local Projects View.

Related Information

[About the information design tool interface](#) [page 22]

[About resources in the information design tool](#) [page 19]




[How to create a universe with the information design tool](#) [page 37]

[About securing resources in the information design tool](#) [page 329]

2.4 About the information design tool interface







Welcome page







The first time you start the information design tool, you see the **Welcome** page. From the **Welcome** page you can access all the resource creation wizards, open existing resources, and link to help and training materials.

You can close the **Welcome** page to display the information design tool interface. To re-open the **Welcome** page, select  **Help**  **Welcome** .

Views and editors

The information design tool interface is composed of views and editors that let you navigate and work on different resources. For more information about each view, see the Related Topics.

View	Description	How to open the view
Local Projects View	Use this view to create and navigate local projects, and to open and validate resources.	Select  Window  Local Projects  .
Repository Resources View	Use this view to navigate repository resources and create secured connections.	Select  Window  Repository Resources  .
Data Foundation Editor	Use this editor to define and maintain the data foundation structure and its connections, and to access the data federation layer.	Double-click a data foundation in the Local Projects View.

View	Description	How to open the view
Business Layer Editor	Use this editor to define and maintain the business layer and its data source.	Double-click a business layer in the Local Projects View.
Connection Editor	Use this editor to edit connection and connection shortcut parameters.	Double-click a connection or connection shortcut in the Local Projects View. To open the editor for a secured connection, double-click the connection in the Repository Resources View.
Project Synchronization View	Use this view to manage shared project resources in the repository and synchronize local resources with the repository.	Select  Window > Project Synchronization  .
Check Integrity Problems	Use this view to review the results of last integrity check.	Select  Window > Check Integrity Problems  .
Query Panel	Use this view to run queries on business layers and published universes.	To run a query on a business layer, in the Business Layer Editor Query pane, edit an existing query, or create a query. To run a query on a published universe, in the Repository Resources View, right-click a universe and select Run Query . You can also run a query on a published universe in the Security Editor.
Security Editor	Use this editor to define security profiles and assign profiles to users.	Select  Window > Security Editor  .

Related Information

[Resetting the user interface display](#) [page 24]
[About local projects and resources](#) [page 78]
[About managing repository resources](#) [page 94]
[About the Data Foundation Editor](#) [page 136]
[About the Business Layer Editor](#) [page 225]
[About the Connection Editor](#) [page 102]
[About project synchronization](#) [page 86]
[Running check integrity](#) [page 316]
[Using the Query Panel](#) [page 297]
[About the Security Editor](#) [page 332]



2.4.1 Resetting the user interface display

Context

The information design tool user interface can be customized by dragging and dropping editor tabs and views, minimizing views, and hiding and splitting panels within the views.

To reset the user interface to the default configuration, select  **Window** > **Reset to Default Display** .

2.5 Setting preferences in the information design tool

To set preferences, from the information design tool main menu, select  **Window** > **Preferences** . The following table gives a brief description of the types of preferences you can set. For more information, see the related topics.

Preference type	Description	
General	These settings and customizations require a good understanding of the Eclipse development environment. To learn more about Eclipse, search for the Eclipse Foundation Web site.	
Help	Lets you select how you would like the help topics to display when you click the help icon. > Content: You can use these settings to include custom help files. These settings require an understanding of the Eclipse help system. To learn more about Eclipse, search for the Eclipse Foundation Web site.	
	Business Layer Editor	Lets you change how object names are generated in relational business layers.
Information Design Tool	Check Integrity	Lets you set the integrity rules to run automatically when saving resources. You can also set the severity level of the rules.
	Data Foundation Editor	Lets you set display options for connections in the data foundation editor. > Appearance: Set display options for columns, tables, and joins in the Data Foundation Editor. > Detections: Set whether or not to automatically detect tables, joins, and cardinalities when inserting tables into the data foundation. Set the default join state for contexts when adding joins and contexts to the data foundation. > Performance: Set options that impact how graphics are displayed in the Data Foundation Editor.
	Languages	Lets you change the language of the user interface and the Preferred Viewing Locale.

Preference type	Description
	Online Tutorials Lets you update the link to the online tutorials.
	Secured Connections Lets you set whether to use the server or the local middleware driver for secured relational connections. <div> <i>i</i> Note This preference applies only if the <i>Download connection locally</i> right is granted in the Central Management Console for the connection. </div>
	Show Values Lets you set how to display table and column values for the Show Values commands.

Related Information

[Setting preferences for the Business Layer Editor](#) [page 25]

[Setting preferences for check integrity](#) [page 26]

[Setting connection display preferences for the Data Foundation Editor](#) [page 27]

[Setting display preferences for the data foundation view](#) [page 28]

[Setting table and join detection options](#) [page 29]

[Setting default join states for contexts](#) [page 31]

[Setting performance-related options for the data foundation view](#) [page 31]

[Setting languages used by the information design tool](#) [page 32]

[Setting a link for online tutorials](#) [page 33]

[Setting middleware for secured relational connections](#) [page 33]




[Setting preferences for showing values](#) [page 34]

2.5.1 Setting preferences for the Business Layer Editor

Context

The Business Layer Editor preference page lets you change how object names are generated in relational business layers.

Procedure

1. From the information design tool main menu, select  **Window**  **Preferences** .

2. In the [Preferences](#) dialog box, expand the **Information Design Tool** node and select **Business Layer Editor**.
3. Select how object names are generated in the business layer:

The option applies to the automatic generation of object names when creating relational business layers and inserting data foundation tables into the business layer .

The option does not apply for business layers based on SAP ERP, SAP NetWeaver BW, and SAP HANA, which use a dedicated strategy for naming object during automatic generation.

Option	Description
Translate table and column names to user-friendly names	If selected, generates object names by changing non-letter characters to blanks and capitalizing the first letter of each word. For example, the column name <code>region_id</code> generates the dimension name <code>Region Id</code> .
Use table and column names as they are	If selected, generates object names using the table and column names in the data foundation. For example, the column name <code>region_id</code> generates the dimension name <code>region_id</code> .

4. To restore the default values for preferences on the current page, click **Restore Defaults**.
5. To save the changes and continue editing preferences, click **Apply**.
6. To save the changes and close the [Preferences](#) dialog box, click **OK**.

Results

The new preferences take effect immediately.

2.5.2 Setting preferences for check integrity

Context

In the check integrity preference page, you can select rules to be run automatically whenever you save a resource. You can also change the severity of the messages returned by each rule.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the [Preferences](#) dialog box, expand the **Information Design Tool** node and select **Check Integrity**.
3. To select rules to be run automatically when you save a resource:
 - a) Select the **Enable background check integrity on save** option.
 - b) Select the rules to include in the background check.

The **Cost** column indicates the relative processing time required to execute the rule.

4. To change the severity of the messages returned by a rule, click in the **Severity** column for the rule. Select a severity from the list.
5. To restore the default values for check integrity preferences, click **Restore Defaults**.
6. To save the changes and continue editing preferences, click **Apply**.
7. To save the changes and close the *Preferences* dialog box, click **OK**.

Results

The background check takes effect immediately.

Related Information

[Running check integrity](#) [page 316]

2.5.3 Setting connection display preferences for the Data Foundation Editor

Context

The Data Foundation Editor preference page lets you to change how connections are displayed in the Data Foundation Editor. For information on other preferences for the Data Foundation Editor, see the Related Topics.

Procedure

1. From the information design tool main menu, select ► **Window** ► **Preferences** ▾.
2. In the *Preferences* dialog box, expand the **Information Design Tool** node and select **Data Foundation Editor**.
3. Set or clear options for how connections are displayed in the data foundation:

Option	Description
Enable automatic expansion	If selected, and the connection provides default qualifiers and/or owners, the catalog in the <i>Connections</i> pane automatically expands the default qualifier/owner.
Show Qualifiers/Owners	If selected, qualifiers and owners, if available, are displayed by default in the <i>Connections</i> pane.

Option	Description
Show only information models (SAP HANA connections)	If selected, in the _SYS_BIC owner of SAP HANA connections, only information models (such as Analytic Views and Calculation Views) are displayed by default in the Connections pane.

- To restore the default values for preferences on the current page, click **Restore Defaults**.
- To save the changes and continue editing preferences, click **Apply**.
- To save the changes and close the [Preferences](#) dialog box, click **OK**.

Results

The new preferences take effect immediately.

Related Information

[Setting display preferences for the data foundation view](#) [page 28]

[Setting table and join detection options](#) [page 29]

[Setting performance-related options for the data foundation view](#) [page 31]

[About connections in the data foundation](#) [page 141]

2.5.4 Setting display preferences for the data foundation view

Procedure

- From the information design tool main menu, select **Window > Preferences**.
- In the [Preferences](#) dialog box, expand the **Information Design Tool** node and then expand the **Data Foundation Editor** node.
- Select **Appearance**.
- Set or clear display options:

The options change the display of elements in the data foundation view.

Option	Description
Show data types	If selected, an icon showing the data type of the column is displayed in front of the column name. For example, AB indicates a string data type, and 1 2 indicates numeric.
Centered	If selected, column names are centered in the table display. Otherwise, column names are left-justified.

Option	Description
Shadow border	If selected, tables display with a shadow border.
Show row counts	If selected, a row count displays for each table.
Postfix alias names with original table name	If selected, the original table name for an alias table is displayed in parentheses after the alias table name.
Prefix table names with owner and qualifier	If selected, the owner and qualifier names are displayed in front of the table name.
Join Lines	Select the type of join line from the list.
Show full table names in join expressions	This option is not currently used.
Auto scroll and zoom to selection	If selected, when you select an element in the view, the view automatically scrolls and zooms to better center the element in the viewing area.

5. To restore the default values for preferences on the current page, click **Restore Defaults**.
6. To save the changes and continue editing preferences, click **Apply**.
7. To save the changes and close the [Preferences](#) dialog box, click **OK**.

Results

The new preferences take effect immediately.

2.5.5 Setting table and join detection options

Context

Set whether or not to automatically detect table keys, row counts, joins, and cardinalities when inserting tables into the data foundation.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the [Preferences](#) dialog box, expand the **Information Design Tool** node and then expand the **Data Foundation Editor** node.
3. Select **Detections**.
4. Set or clear detection options:

Option	Description
Detect keys	If selected, when inserting a table into the data foundation, automatically detects primary and foreign keys in the data source and sets the keys in the data foundation table.
Keep keys defined in data foundation if none detected in database	If selected, when you use the Detect Keys command, if no keys are detected for a table in the data source, the keys that were set manually in the data foundation table are kept.
Detect row counts	If selected, when inserting a table into the data foundation, the number of rows in the table is counted and stored in the data foundation.
Detect joins	If selected, when inserting tables into the data foundation, joins are detected automatically (using the selected method) and inserted into the data foundation.
Detect cardinalities	<p>If selected, when inserting tables into the data foundation, the cardinality of joins is detected and set automatically.</p> <div> <p>i Note</p> <p>Detect joins must also be selected.</p> </div> <p>For a description of the method used to detect cardinality, see the related topic about cardinality.</p>

5. To restore the default values for preferences on the current page, click **Restore Defaults**.
6. To save the changes and continue editing preferences, click **Apply**.
7. To save the changes and close the *Preferences* dialog box, click **OK**.

Results

The new preferences take effect immediately.

Related Information

[About table keys](#) [page 157]

[About table row counts](#) [page 159]

[Detecting joins](#) [page 163]

[About cardinality](#) [page 165]

[Setting default join states for contexts](#) [page 31]

2.5.6 Setting default join states for contexts

Context

Application preferences let you set the default join state for contexts when adding joins and contexts to the data foundation.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the **Preferences** dialog box, expand the **Information Design Tool** node and then expand the **Data Foundation Editor** node.
3. Select **Detections**.
4. Select the state that new joins will have when they are added to existing contexts:

Option	Description
Neutral	This is the default value. Joins added to the data foundation are neutral in any existing contexts. Neutral joins are not explicitly included or excluded but may be used in a query path.
Excluded	Joins added to the data foundation will be excluded in any existing contexts.
Included	Joins added to the data foundation will be included in any existing contexts.

5. Optionally, select the check box to apply this rule when creating contexts.
By default, when creating a context, all joins in the data foundation are neutral. If, for example, you set the default state for joins to **Excluded** in step 4, and if you select **Also apply this rule when creating contexts**, then when creating a context, all joins in the data foundation will be excluded from the context.

Related Information

[About contexts](#) [page 175]

2.5.7 Setting performance-related options for the data foundation view

Context

The following options enhance the display in the data foundation view. In some cases the enhancements can cause sluggishness when dragging elements within the view. Options can be unselected in cases where this results in a better display-performance benefit.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the *Preferences* dialog box, expand the **Information Design Tool** node and then expand the **Data Foundation Editor** node.
3. Select **Performance**.
4. Set or clear performance options:

Option	Description
Use transparency effects	If selected, when dragging a table in the data foundation view, a semi-transparent shadow of the table follows to show the trail.
Use line smoothing	If selected, displays smoothed lines for joins.
Use image enhancement when zooming	If selected, avoids using big pixels when zooming.
Use text line smoothing	If selected, lines in the text are smoothed.
Use fading transitions	This option is not currently used.

5. To restore the default values for preferences on the current page, click **Restore Defaults**.
6. To save the changes and continue editing preferences, click **Apply**.
7. To save the changes and close the *Preferences* dialog box, click **OK**.

Results

The new preferences take effect immediately.

2.5.8 Setting languages used by the information design tool

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the *Preferences* dialog box, expand the **Information Design Tool** node and select **Languages**.
3. To change the language of the user interface, select the language from the *Product Languages* list.
4. To change the **Preferred Viewing Locale**, select the language from the list.
For information about the Preferred Viewing Locale and how it impacts the language display, see the related topic about multilingual universes.
5. To save the changes and continue editing preferences, click **Apply**.
6. To save the changes and close the *Preferences* dialog box, click **OK**.
7. Exit and restart the information design tool for the language change to take effect.

Related Information

[Multilingual universes](#) [page 53]

2.5.9 Setting a link for online tutorials

Context

You can access online tutorials about the information design tool from the **Help** menu. The [Online Tutorials](#) preferences page allows you to update the URL address to the tutorials.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the [Preferences](#) dialog box, expand the **Information Design Tool** node and select **Online Tutorials**.
3. Enter the new URL address in **Online Tutorial Address**.
4. To save the change and continue editing preferences, click **Apply**.
5. To save the change and close the [Preferences](#) dialog box, click **OK**.

Results

The new address takes effect immediately.

Related Information

[How to get help using the information design tool](#) [page 35]

2.5.10 Setting middleware for secured relational connections

Prerequisites

The secure connections middleware preference applies only if the [Download connection locally](#) right is granted in the Central Management Console for the connection.

Context

When running queries on secured relational connections in the information design tool, you can choose to run queries on the server using the server middleware driver, or locally using the local middleware driver.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the *Preferences* dialog box, expand the **Information Design Tool** node and select **Secured Connections**.
3. Select the middleware to use:

Option	Description
Server middleware	To use the middleware driver on the repository server.
Local middleware	To use the middleware driver on the local machine.

4. To restore the default value, click **Restore Defaults**.
5. To save the changes and continue editing preferences, click **Apply**.
6. To save the changes and close the *Preferences* dialog box, click **OK**.

Related Information

[About secured connections](#) [page 100]

2.5.11 Setting preferences for showing values

Context

For the show values commands in the data foundation and business layer editors, you can select how to display values.

Procedure

1. From the information design tool main menu, select **Window > Preferences**.
2. In the *Preferences* dialog box, expand the **Information Design Tool** node and select **Show Values**.
3. Select how you want the values to display.
4. To save the change and continue editing preferences, click **Apply**.
5. To save the change and close the *Preferences* dialog box, click **OK**.

Results



The new display option takes effect immediately.

Related Information

[Showing values in a data source](#) [page 185]

2.6 How to get help using the information design tool

The **Help** menu in the information design tool lets you link to different types of user assistance for the application.

Help menu command	Description
Welcome	From the Welcome page you can access all the resource creation wizards, open existing resources, and link to help and training materials.
Cheat Sheets	<p>A cheat sheet shows how to complete a complex task, for example, creating a relational universe.</p> <p>The Cheat Sheets command displays a list of available cheat sheets. Double-click a cheat sheet to start. It opens in the information design tool Help View.</p> <p>For some steps, you can click Click to perform and the application will start the appropriate wizard to help you complete that step.</p> <p>To get more help on a step, click the help icon .</p>
Online Tutorials	<p>The Online Tutorials command links to the list of official product tutorials for the information design tool on the SAP Community Network.</p> <p>If the URL address of the online tutorials changes, you can enter the new address in the information design tool preferences.</p>
Help Contents	<p>The Help Contents command opens the <i>Information Design Tool User Guide</i> in a help window. To display help topics, you can navigate the Table of Contents, search the text, or look up topics in the index.</p> <div><p> Note</p><p>See the SAP Help Portal for the most updated version of the guide at http://help.sap.com/.</p></div>

Help menu command	Description
Search	The Search command opens the Help View on the search function. To search the contents of the <i>Information Design Tool User Guide</i> , enter text in the Search expression field.

Related Information

[Setting a link for online tutorials](#) [page 33]

2.7 Wizards in the information design tool

You can use wizards to help you create local resources in the information design tool. The wizards are available from the **New** menu in the main tool bar. To see a list of all wizards, select **New > Other**. To start a wizard, select it in the list and click **Next**.

To get help on a particular page of any wizard, click the help icon in the wizard dialog box.

Related Information

[About resources in the information design tool](#) [page 19]

[How to create a universe with the information design tool](#) [page 37]

3 Creating Universes

3.1 How to create a universe with the information design tool

Prerequisites

Before you begin:

- Make sure the middleware drivers are configured for the data sources to which you want to connect. For more information about middleware configuration, see the *Data Access Guide*. For information about supported data sources, see the SAP Business Objects BI Platform 4.1 Supported Platforms (PAM) at <http://service.sap.com/pam>.
- Make sure you have the appropriate rights defined in the Central Management Console (CMC). See the related topic about CMC rights for information design tool users.
- Decide if the data foundation type should be single-source or multisource-enabled. The type and number of connections available, as well as the SQL syntax that is used to define SQL structures depends on the data foundation type. For more information, see the related topic about data foundation types. Connections for multisource-enabled data foundations must be secured, relational connections, and are managed by the data federation service. Information about tuning the data federation service can be found in the *Data Federation Administration Tool Guide*.
- Refer to the additional information available if you are creating a universe on one of the following data sources:
 - [Using SAP NetWeaver BW data sources](#) [page 40]
 - [Using SAP HANA data sources](#) [page 45]
 - [Using SAP ERP data sources](#) [page 48]
 - [Using Microsoft Analysis Services \(MSAS\) data sources](#) [page 49]
 - [Using Essbase data sources](#) [page 51]
 - [Using SAS data sources](#) [page 53]
 - [Multilingual universes](#) [page 53]

Note

For OLAP universes, you do not need to create a data foundation. The business layer is built directly from the objects you select in the source cube.

Context

You can use the [New Universe](#) wizard to create the resources you need to publish a local universe: either a single-source relational or OLAP universe. If you want to base your universe on secured connections, the connection shortcuts must exist in a local project. The wizard lets you create only local connections. For more information, see the related topic. To start the wizard, select **File > New Universe**.

The following procedure describes how to make any kind of universe from scratch. Links to more information on each step in the procedure can be found in the Related Topics.

Procedure

1. Create a local project. In the Local Projects View select **File > New > Project**.

The resources you use to build the universe are created and stored in the project

2. Define the connections. Connections can be local or secured:
 - Use a local connection if you want to publish the universe on the local file system. Later, you can publish the business layer to a repository.
 - Create a secured connection if you want to create a multisource-enabled universe, or if you want to publish the universe in a repository without first publishing locally. For secured connections, you must create connection shortcuts in the local project to reference the secured connections in the repository.

Option	Command
To create a local relational connection	In the Local Projects View, right-click the project folder and select New > Relational Connection .
To create a local OLAP connection	In the Local Projects View, right-click the project folder and select New > OLAP Connection .
To create a secured relational connection	<p>In the Repository Resources View, start a repository session. Right-click the Connections folder or sub-folder and select Insert Relational Connection.</p> <p>To create a connection shortcut, select the connection in the Connections folder of the repository and select Create Relational Connection Shortcut.</p>
To create a secured OLAP connection	<p>In the Repository Resources View, start a repository session. Right-click the Connections folder or sub-folder and select Insert OLAP Connection.</p> <p>To create a connection shortcut, select the connection in the Connections folder of the repository and select Create OLAP Connection Shortcut.</p>

3. Create the data foundation (relational data sources only). In the Local Projects View, right-click the project folder and select **New > Data Foundation**.

- For a single-source universe, select a single connection identifying the database source.
- To build a data foundation with multiple relational connections, create a multisource-enabled data foundation.

The data foundation opens in the editor. To build the structure of the data foundation, see the related topic.

4. Create the business layer. In the Local Projects View, right-click the project folder and select **New > Business Layer**.

- For relational business layers, select the data foundation to be the basis for the business layer. You can choose to automatically generate objects in the business layer for all data foundation structures, or select columns to be mapped as objects.





Note



To build a universe on multiple data sources (relational data sources only), the business layer must be based on a multisource-enabled data foundation.

- For OLAP business layers, select the connection to the OLAP cube. Objects are created automatically for all structures in the cube.

The business layer opens in the editor. To build the business layer, see the related topic.

5. In the business layer, you can create and run queries to validate and test the universe.
6. Publish the business layer:
 - Business layers based on local connections must be published to a folder on the local file system. Later, you can publish the resulting local universe to a repository. See the related topic about publishing a local universe to the repository.
 - Business layers based on one or more secured connections must be published to the repository on the same Central Management System where the secured connections are stored.

Option	Command
To publish the universe locally	In the Local Projects View, right-click the business layer and select  Publish > To a Local Folder  .
To publish the universe to a repository	In the Local Projects View, right-click the business layer and select  Publish > To a Repository  .

7. Define universe security. To open the Security Editor, on the information design tool main menu, select  **Window** > **Security Editor** . Open a session on the repository where the universe is published.
Use the Security Editor to define security profiles on the published universe. You also use the Security Editor to assign profiles to users and groups.

Related Information

[About the information design tool interface](#) [page 22]
[About resources in the information design tool](#) [page 19]
[CMC rights for information design tool users](#) [page 330]
[About data foundation types](#) [page 133]
[Creating a local project](#) [page 79]
[Creating a relational connection](#) [page 102]
[Creating an OLAP connection](#) [page 123]
[Creating a connection shortcut](#) [page 129]
[How to build a data foundation](#) [page 138]
[How to build a relational business layer](#) [page 218]
[How to build an OLAP business layer](#) [page 222]
[About queries in a business layer](#) [page 290]
[Publishing a universe](#) [page 324]
[Publishing a local universe to the repository](#) [page 326]

3.2 Using SAP NetWeaver BW data sources

Direct access to BEx Queries

SAP BusinessObjects query and reporting applications can use direct access to access data in a single BEx Query. You do not need to build a universe. Define an OLAP connection to SAP NetWeaver BW that uses the **SAP BICS Client** middleware driver. When defining the connection, select the option to specify a cube in the connection and select the BEx Query.


Universes on SAP NetWeaver BW

To build a universe on SAP NetWeaver BW, you must create a multisource-enabled data foundation based on a secured relational connection to SAP NetWeaver BW. You then build the business layer on this data foundation. For detailed steps on how to build a universe, see the related topic.

i Note

For information on the authorizations needed to allow users of query and reporting applications to access multisource-enabled universes on SAP NetWeaver BW, see SAP Note #1465871.

Relational connections to SAP NetWeaver BW are managed by the data federation service. For information about optimizing queries, see the *Data Federation Administration Tool Guide*.

To see the InfoProviders supported for relational connections to SAP NetWeaver BW, see the Data Access for the Semantic Layer section of the SAP Business Objects BI Platform 4.1 Supported Platforms (PAM) at <http://service.sap.com/pam> .

When you add an SAP NetWeaver BW connection to a data foundation, by default tables and joins are automatically inserted. When you create a business layer on the data foundation, by default objects are automatically inserted into the business layer.

To turn off automatic insertion, unselect the **Detect tables** option in the advanced properties of the connection when you are adding the connection into the data foundation. To turn off automatic insertion of business layer objects, unselect the **Automatically create folders and objects** option when selecting the data foundation in the New Business Layer wizard.

See the related topic for information on how the objects in an InfoProvider map to the objects that are automatically inserted into the data foundation and business layer in the information design tool.

Refreshing universes based on SAP NetWeaver BW

When objects are added to the underlying InfoProvider, several commands exist to help you update the data foundation and business layer with the changes. The recommended procedure is described in the related topic

Related Information

[Creating an OLAP connection](#) [page 123]

[How to create a universe with the information design tool](#) [page 37]

[How InfoProvider objects are mapped in a universe](#) [page 41]



[Refreshing universes based on SAP Netweaver BW](#) [page 44]

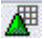






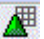

3.2.1 How InfoProvider objects are mapped in a universe






When you add an SAP NetWeaver BW connection to a data foundation, by default tables and joins are automatically inserted. When you create a business layer on the data foundation, by default objects are automatically inserted into the business layer.

The following table describes how the objects in an InfoProvider map to the objects that are automatically inserted into the data foundation and business layer in the information design tool. The naming convention for the automatically generated data foundation tables are as follows:

- I-Table: A table with a name prefixed by I (InfoCube) that maps to the InfoProvider fact table.
- D-Table: A table with a name prefixed by D (Dimension) that maps to an InfoProvider master data table.
- T-Table: A table with a name prefixed by T (Text) that maps to an InfoProvider text table.

InfoProvider object	Data foundation	Business layer
Fact table	<ul style="list-style-type: none">• Inserts an I-Table mapped to the fact table.• Inserts D-Tables mapped to each master data table.• Inserts T-Tables mapped to each text table.	
Dimension 		 Inserts a folder in the business layer for each InfoProvider dimension (except for the Data Package and Unit dimensions). Business layer objects for the characteristics in the Unit dimension are inserted into the folder for the re-

InfoProvider object	Data foundation	Business layer
		lated key figure. The Data Package dimension is not mapped.
Characteristic 	Inserts a column in the data foundation fact table that maps to the master data value.	 Inserts a dimension object in the dimension folder associated with the I-Table column.
	<p>Creates an alias of the associated T-Table and links it to the column in the I-table.</p> <p>The T-Table contains columns for the description of the master data in short, medium, or long format. The table also has a column CAPTION which contains the longest description available for the characteristic.</p> <div> <p>i Note</p> <p>A characteristic of data type DATS or TIMS has no associated text table because date and time characteristics do not have descriptions.</p> </div>	 Inserts attribute objects under the dimension for each column in the T-Table.
Display attribute* 	If a characteristic contains at least one display attribute, inserts an alias of the D-Table. Inserts a column in this table mapped to the display attribute. The D-Table contains a column for each display attribute of the characteristic.	 Inserts a folder under the dimension named for the parent characteristic.  Inserts dimension objects in this folder for each column in the D-Table.
	Inserts an alias of the T-Table for each display attribute.	 Inserts attribute objects under the dimension object for each column in the T-Table.
Navigation attribute* 	<p>Inserts a column in the I-Table and a column in the D-table mapped to the navigation attribute.</p> <p>T-Tables are linked directly to I-Table for both the parent characteristic and its navigation attributes.</p>	 Inserts a dimension object in the dimension folder associated with the I-Table column. <p>The dimension object for the navigation attribute is at the same level as the dimension for the parent characteristic but not necessarily adjacent.</p>

InfoProvider object	Data foundation	Business layer
		<p>➔ Tip</p> <p>Filters on the object for the navigation attribute are more efficient than filters on the display attribute object. When filtering a navigation attribute, the fact table is filtered directly.</p>
Key figure 	Inserts a column in the I-Table with the technical name of the key figure.	<p> For a key figure without unit or currency, inserts a measure in the Measures folder.</p> <p> For a key figure with unit or currency, inserts a sub-folder in the Measures folder.</p> <p> Inserts a measure in the sub-folder for the key figure.</p> <p> Inserts a dimension in the sub-folder for each unit or currency characteristic.</p>
Time-dependent data	<p>If the InfoProvider contains time-dependent data, creates input columns in the appropriate data foundation tables to handle time-dependent data.</p> <p>Creates a parameter in the data foundation called key date for each input column. By default, at query time, the key date parameter is not prompted. It is automatically assigned the current date. You can change this behavior by editing the key date parameter. For more information on input columns and editing parameters, see the related topics.</p>	The business layer inherits the key date parameter.

* If the flag “Attribute only” is checked for a display or navigation attribute in the InfoProvider, it is not exposed in the data foundation.

Related Information

[Using SAP NetWeaver BW data sources](#) [page 40]

[About business layer objects](#) [page 215]

[About input columns in the data foundation](#) [page 178]

[Inserting and editing a parameter](#) [page 281]

3.2.2 Refreshing universes based on SAP Netweaver BW

Context

Use this procedure when objects change in the InfoProvider of an SAP NetWeaver BW data source and you want to reflect the changes in the universe. You can link to more information about each step in the Related Topics.

Procedure

1. Refresh the structure of the data foundation.

Refreshing the data foundation compares the existing tables in the data foundation with those in data source and proposes updates to the data foundation tables: deletes obsolete tables and columns, inserts missing columns, and updates changed columns.

2. Synchronize the tables in the data foundation.

Synchronizing tables searches the data source for new tables (using the SAP NetWeaver BW strategy), and inserts the new tables and joins into the data foundation.

3. Save the data foundation.

4. Refresh the business layer by inserting candidate objects.

Inserting candidate objects searches the data source for new objects (using the SAP NetWeaver BW strategy) and updates the business layer.

Inserting candidate objects does not detect obsolete objects in the business layer. You must find and delete obsolete objects manually.

Results

Inserting candidate objects updates the business layer independently of the data foundation. If you do not also refresh the structure and synchronize tables in the data foundation, you can introduce inconsistencies between the data foundation and business layer.

Related Information

[About refreshing a data foundation](#) [page 188]

[Synchronizing tables](#) [page 189]

[Inserting candidate objects](#) [page 293]

[Using SAP NetWeaver BW data sources](#) [page 40]

3.3 Using SAP HANA data sources

Direct access to SAP HANA information models

SAP Crystal Reports for Enterprise can use direct access to access data in a single information model (such as an Analytic View or Calculation View). You do not need to create a universe. Define an OLAP connection to SAP HANA that uses the **SAP HANA Client** middleware driver. In the connection, you can choose to not specify a cube. In this case, the user is prompted to select an information model at query run time. You can also specify an information model in the connection so that all queries are associated with that information model.

Building universes on SAP HANA

Universes on SAP HANA are based on relational connections. There are two ways to create resources to build a universe on SAP HANA:

- Use the New Data Foundation and New Business Layer wizards to create a data foundation and business layer separately.

This method allows you to include tables in the data foundation. You can also include information models.

Note

If your data foundation will include only tables, the connection to SAP HANA can use an ODBC middleware driver. As soon as you include an information model in the data foundation, the connection to SAP HANA must use a JDBC middleware driver.

Create a single-source data foundation unless you need to federate data from multiple sources. A multisource-enabled data foundation is necessary if you need to access data in more than one SAP HANA server or instance.

Select the tables or information models that you want to include in the data foundation. Any columns hidden in an SAP HANA view are also hidden in the data foundation table. It is not recommended to create joins between tables representing SAP HANA views as this can impact performance.

The New Business Layer wizard automatically creates the dimensions and attributes in each information model in a business layer folder. For single-source data foundations, the wizard uses metadata from the SAP HANA view to create measures in the business layer with the appropriate aggregation function.

For more information, see the related topic on how to build a universe.

- Use the New SAP HANA Business Layer wizard to automatically create a single-source data foundation and business layer based on selected SAP HANA information models.

Note

The connection to SAP HANA must use a JDBC middleware driver.

This wizard creates the data foundation tables for the selected SAP HANA views. Any columns hidden in an SAP HANA view are also hidden in the data foundation table. The wizard then creates the dimensions and measures (with appropriate aggregation functions) defined in the SAP HANA views.

The advantage of this method is that for any dimensions and attributes that are common to different views, the wizard creates a single business layer object.

The wizard also creates contexts in the data foundation and aggregate awareness in the business layer so that when you run a query on the universe, the experience is the same as accessing a standard relational universe:

- For all measures and dimensions common to the SAP HANA views accessed in the query, results display in a single block.
- For dimensions that are not common to the SAP HANA views accessed in the query, results display in separate blocks.

For more information, see the related topic on creating relational resources on SAP HANA information models.

Browsing SAP HANA connections

When browsing the connection in the data foundation, different views are identified by table types each with its own icon. You can filter the tables in the connection by table type. There is also a filter, set by default, to show only tables representing information models. For more information, see the related topics about filtering tables in the connection.

Note

An Analytic View sometimes appears as a Calculated View table type in the connection. This happens when the Analytic View contains a calculated measure in the SAP HANA model.

Showing values in a data foundation based on SAP HANA

In single-source data foundations, when showing table and column values for a table corresponding to an Analytic View, the information design tool aggregates the values in columns that represent measures by using the aggregation function defined in the model. The measure values are grouped by the selected columns that represent attributes. For example, if you show values for **Product** and **Sales** in an Analytic View where **Sales** is aggregated with the sum function, the results show **Sales** by **Product**.

SAP HANA variables in the data foundation

In single-source data foundations, variables and input parameters in SAP HANA information models are associated with the corresponding tables in the data foundation.

When showing values in the data foundation or running a query in the Query Panel, you are prompted to enter values for the variables and parameters.

You can get information about variables and parameters in the data foundation properties **Variables** tab. Also, when refreshing the structure of the data foundation, any variables in the view that are added, deleted, or modified are taken into account.

You may want to enhance the data foundation by creating derived objects: derived tables, calculated columns, or custom SQL lists of values. Since the SAP HANA variables are hidden, you need to manage the variables directly in the SQL expressions for the objects that you create. For recommendations and limitations on enhancing data foundations that contain SAP HANA variables, see SAP Note 1913504.

If any SAP HANA variables or input parameters change in the underlying SAP HANA information model, you must run a refresh structure on the data foundation in the information design tool.

Related Information

[Creating an OLAP connection](#) [page 123]

[How to create a universe with the information design tool](#) [page 37]

[Creating relational resources on SAP HANA information models](#) [page 47]

[Filtering the tables in the connection by table type](#) [page 147]

[Filtering tables in an SAP HANA connection by information model](#) [page 148]

[Showing SAP HANA variable information](#) [page 184]

[About refreshing a data foundation](#) [page 188]

3.3.1 Creating relational resources on SAP HANA information models

Context

The New HANA Business Layer wizard automatically creates a data foundation and business layer based on selected SAP HANA information models.

The wizard creates a single-source data foundation referring to the local connection to SAP HANA that you provide. The data foundation contains a table for each view. The tables are not joined.

The wizard creates a business layer that contains the dimensions and measures defined in the SAP HANA views. Dimensions and attributes that are common to different views are bound to a single business layer object.

Before you begin, you need the following resources:

- A local project
- In the local project, a local connection to SAP HANA.

Note

The connection must use a JDBC middleware driver.

Procedure

1. In the Local Projects View, right-click the project and select **New > SAP HANA Business Layer**.
2. Enter a name for the business layer and data foundation.
By default, the data foundation has the same name. You can change the data foundation name.
3. Optionally enter a description for the business layer and click **Next**.
4. Select a connection and click **Next**.
Only local, relational connections to SAP HANA in the local project are listed.
5. Select one or more activated views to be the basis of the business layer and click **Finish**.
A data foundation and business layer are created in the local project. The business layer opens in the editor.

Next Steps

You can publish the business layer to a local folder. If you want to publish the universe to a repository, see the related topic.

If any SAP HANA variables or input parameters change in the underlying SAP HANA information model, you must run a refresh structure on the data foundation in the information design tool.

Related Information

[Publishing a local connection to the repository](#) [page 326]

[Changing a connection in a data foundation](#) [page 143]

[Publishing a universe](#) [page 324]

[Publishing a local universe to the repository](#) [page 326]

[Using SAP HANA data sources](#) [page 45]

[About refreshing a data foundation](#) [page 188]

3.4 Using SAP ERP data sources

To build a universe on SAP ERP, you must create a data foundation on a relational ERP connection. Then build the business layer on this data foundation. For detailed steps on building a universe, see the related topic.

When you create a relational connection to the SAP ERP data source, the InfoSets, SAP Queries, and ABAP functions in the data source are exposed as tables in the connection. For more information about how ERP data sources are mapped in the connection, see the *Data Access Guide*.

The data foundation can be single-source to support local connections. Single-source data foundations support joins between tables, with the following restrictions:

- You can detect joins only based on database keys in the ERP data source. You must first detect database keys in the data foundation.
- You cannot insert joins manually, insert calculated columns, or insert column filters.

For support of calculated columns, filters, and manual joins, create a multisource-enabled data foundation on a secured connection.

When you insert a table into the data foundation, the table type of InfoSet, SAP Query, or ABAP Function is saved as a table property in the data foundation.

When you insert an ABAP Function table, one data foundation table is created to map the main function. The table contains input columns for the input parameters of the function. These parameters can be mandatory or optional. To assign a value to mandatory parameters, you need to edit the input columns. To do this, see the related topic.

When creating the business layer, the object names are automatically generated from the column descriptions in the data foundation, rather than from the column names. For reference, the column names are saved as the description of the business layer object.

Restriction

Measures containing aggregate functions cannot be used as filters in the Query Panel. This limitation is due to the fact that the resulting SQL expression contains the HAVING clause, which is not supported by the SAP ERP connection. If you add a measure containing an aggregation function as a filter, an error occurs when you refresh the query.

Related Information

[How to create a universe with the information design tool](#) [page 37]

[Editing input columns](#) [page 179]

[About parameters and lists of values in the data foundation](#) [page 180]

[About multisource-enabled data foundations](#) [page 134]

3.5 Using Microsoft Analysis Services (MSAS) data sources

When you create a business layer on an MSAS data source, the business layer objects are generated automatically.

The business layer can be refreshed to reflect changes in the underlying cube using the **Refresh Structure** command in the Business Layer Editor **Actions** menu.

The table below gives details about how certain objects in the MSAS cube are mapped in the business layer.

MSAS object	Business layer mapping
Perspective	When you create a business layer, in the Select OLAP Connection page of the New Business Layer wizard, the base cube in the MSAS data source is listed

MSAS object	Business layer mapping				
	<p>first in the list of connection cubes. Additional cubes and perspectives in the data source are mapped as cubes and are listed in alphabetical order.</p> <p>The cube you select in the list of connection cubes becomes the basis for the objects in the business layer.</p>				
Dimension	Analysis dimensions are created in the business layer for each dimension in the cube.				
Display folder	Folders are created in the analysis dimension to contain the hierarchies in the display folder.				
Hierarchy	<p>For value-based (parent-child) hierarchies, a value-based hierarchy is created in the analysis dimension. The attributes are created in the Attributes folder in the hierarchy.</p> <div> <p>i Note</p> <p>Unbalanced hierarchies are supported.</p> </div> <p>For level-based hierarchies, a business layer dimension is created in the analysis dimension. A hierarchy is created in the business layer dimension with the levels and their properties (as level attributes) in the Levels folder.</p>				
Attribute hierarchy	Attribute hierarchies in the cube are created as level-based hierarchies in the analysis dimension.				
Named set	Named sets are created in the related analysis dimension, in the folder Named sets .				
Measure group	Folders are created to contain the measures in measure groups and sub-groups.				
Measure Calculated measure	Measures and calculated measures are created as measures in the appropriate measure group folder. A measure attribute is created for the formatted value.				
KPI	<p>KPIs are not exposed in connection metadata, but you can make the KPI values available to queries by creating measures in the business layer with the MDX functions KPIValue and KPIGoal.</p> <p>For example, if the cube contains a KPI with the name Operating Profit, you can create measures in the business layer with the following MDX expressions. In this example, the measures are created in the measure group folder \subfolder Performance\Profit. The name of the KPI value in the cube is surrounded by double quotes in the MDX function.</p> <table> <tr> <th>Measure name</th><th>MDX Expression</th></tr> <tr> <td>Actual Profit</td><td>KPIValue("Operating Profit")</td></tr> </table>	Measure name	MDX Expression	Actual Profit	KPIValue("Operating Profit")
Measure name	MDX Expression				
Actual Profit	KPIValue("Operating Profit")				

MSAS object	Business layer mapping								
	<table> <tr> <th>Measure name</th><th>MDX Expression</th></tr> <tr> <td>Profit Target</td><td>KPIGoal("Operating Profit")</td></tr> <tr> <td>Profit Variance</td><td>(@Select(Performance\Profit\Actual Profit) - @Select(Performance\Profit\Profit Target)) / abs(@Select(Performance\Profit\Profit Target))</td></tr> <tr> <td>Profit Pct Achieved</td><td>IIF(ISEMPTY(@Select(Performance\Profit\Profit Target)), null, @Select(Performance\Profit\Profit Variance) +1)</td></tr> </table>	Measure name	MDX Expression	Profit Target	KPIGoal("Operating Profit")	Profit Variance	(@Select(Performance\Profit\Actual Profit) - @Select(Performance\Profit\Profit Target)) / abs(@Select(Performance\Profit\Profit Target))	Profit Pct Achieved	IIF(ISEMPTY(@Select(Performance\Profit\Profit Target)), null, @Select(Performance\Profit\Profit Variance) +1)
Measure name	MDX Expression								
Profit Target	KPIGoal("Operating Profit")								
Profit Variance	(@Select(Performance\Profit\Actual Profit) - @Select(Performance\Profit\Profit Target)) / abs(@Select(Performance\Profit\Profit Target))								
Profit Pct Achieved	IIF(ISEMPTY(@Select(Performance\Profit\Profit Target)), null, @Select(Performance\Profit\Profit Variance) +1)								
Visual totals	You can use the MSAS-specific function VisualTotals in the MDX expressions for business layer objects.								

Related Information

[How to create a universe with the information design tool](#) [page 37]

[OLAP data source properties](#) [page 228]

[About business layer objects](#) [page 215]

[Inserting and editing measures](#) [page 249]

[Refreshing an OLAP business layer](#) [page 293]

3.6 Using Essbase data sources

When you create a business layer on an Essbase data source, the business layer objects are generated automatically.

The business layer can be refreshed to reflect changes in the underlying cube using the **Refresh Structure** command in the Business Layer Editor **Actions** menu.

Restriction

When you insert an MDX object into an Essbase business layer (for example, a named set, calculated member, or measure), be sure that the object name is not the same as data in the cube. For example, if Region is the name of a hierarchical level in the cube, you cannot name a new MDX object Region. If the new object is given the same name as data in the cube, the object is not usable in a query.

The table below gives details about how certain objects in the Essbase cube are mapped in the business layer.

Essbase object	Business layer mapping
Dimension	Analysis dimensions are created in the business layer for each dimension in the cube.
Accounts dimension	<p>In the New Business Layer wizard, you select which dimension to use to create the measures in the business layer. By default this is the dimension flagged as the Accounts dimension. Measures are created in the business layer for each object in the dimension. The organization of the measures in the Essbase outline is maintained in the business layer.</p> <div> <p>➔ Tip</p> <p>For certain applications you may want to specify a dimension other than the Accounts-type dimension for the measures. In this case, the measures are created as an analysis dimension in the business layer and the analysis capabilities of hierarchies, such as member selection, are available.</p> </div>
Hierarchy	<p>For each hierarchy in the cube, a hierarchy is created in the analysis dimension. All hierarchies are generated as value-based.</p> <div> <p>i Note</p> <p>When opening the business layer in the Query Panel, hierarchy levels are determined spontaneously and can be selected in the query. Also, you can insert levels into the business layer.</p> </div>
User Defined Attribute (UDA)	UDAs are created as named sets defined on the associated hierarchy and appear in the analysis dimension.
Attribute	Attributes are created in an Attributes folder in the hierarchy.
Attribute hierarchy	If the attributes are designed as a hierarchy in the cube, an attribute hierarchy is also created in the analysis dimension.
Dynamic Time Series (DTS)	A DTS is not generated automatically in the business layer, but you can use MDX functions, such as HTD, QTD (history-to-date, quarter-to-date) in the object definitions.
Substitution variable	<p>Substitution variables are not exposed in the business layer, but you can invoke a substitution variable in an MDX expression. The substitution variable name must be prefixed with the ampersand (&) character.</p> <p>For example, if the cube contains a variable named CurrentMonth, you can use the variable in the definition of a named set:</p> <pre>WITH SET [Current Month] AS '([Time].[&CurrentMonth])'</pre> <p>Example of the substitution variable in the definition of a calculated member:</p> <pre>WITH MEMBER [Measures].[Current Month Quantity] AS '([Measures].[Quantity Sold], [Time].[&CurrentMonth])'</pre>

Related Information

[How to create a universe with the information design tool](#) [page 37]

[About business layer objects](#) [page 215]

[Refreshing an OLAP business layer](#) [page 293]

3.7 Using SAS data sources

To build a universe on SAS, you must create a multisource-enabled data foundation on a secured connection. Then build the business layer on this data foundation. For detailed steps on building a universe, see the related topic.

Connections to SAS are managed by the data federation service. For information about optimizing queries to SAS data sources, see the *Data Federation Administration Tool Guide*.

Related Information

[How to create a universe with the information design tool](#) [page 37]

3.8 Multilingual universes

The information design tool supports the creation of multilingual universes. This feature enables a multilingual solution using a single universe metadata model:

- The designer creates the universe in the source language in the information design tool.
- Translators translate the metadata in the data foundation and business layer using the translation management tool. For more information about translating metadata, see the related topic.
- Report designers can then build reports once from the same universe that can be displayed in several languages based on user preferences.

Three language parameters impact how labels, metadata, and data are displayed in the information design tool:

- The Product Language determines the language of the user interface of the information design tool. This parameter is set in the information design tool preferences.
- The Preferred Viewing Locale is the user's preferred language for viewing report and query objects in an application. This parameter is set in the information design tool preferences.
A locale defines a language and a geographical area. Locale abbreviations consist of the language abbreviation followed by the country abbreviation, for example, fr_FR. A locale also defines the way data is sorted and how dates and numbers are formatted. Data is displayed in a fallback locale when viewing a translated document and no translation in the user's Preferred Viewing Locale is available. The fallback locale can be defined (in the translation management tool), or defaults to the dominant locale that is automatically defined for every locale.

-
- **Connection Language:** For data sources that support a language parameter, the language parameter is entered when you create or edit a connection. This determines the language of the data.

Metadata source language in the information design tool

The metadata of the data foundation (table and column names) is created in the language of the metadata in the data source. The metadata that you insert into the data foundation can be entered in any language.

For SAP NetWeaver BW connections, the data foundation can be generated automatically in the language specified in the connection language parameter.

The metadata of a relational business layer is created in the language of the data foundation metadata. For OLAP business layers, the metadata is created in the language of the connection language parameter. The metadata that you insert into the business layer can be entered in any language.

Once the business layer is generated, the metadata (when viewed in the Business Layer Editor) remains in that language, even if you change the connection language parameter.

When designing the universe, with the `@Variable` function, you can use the `PREFERRED_VIEWING_LOCALE` and the `DOMINANT_PREFERRED_VIEWING_LOCALE` variables to customize the universe in order to filter multilingual data and retrieve only data in the user's Preferred Viewing Locale at query time.

Multilingual display in the information design tool

The Preferred Viewing Locale (defined in the information design tool preferences) determines the language of the metadata and data in the Query Panel, as long as two conditions are met:

- Translations are available in the language (metadata)
- The connection supports the language parameter (data)

When displaying values from the data source in the Connection Editor, the metadata and data displays in the language of the data source according to the current value of the connection language parameter.

Related Information

[Translating universe metadata](#) [page 55]

[How to create a universe with the information design tool](#) [page 37]

[Setting languages used by the information design tool](#) [page 32]

3.8.1 Translating universe metadata

Prerequisites

Before you begin, the universe you want to translate must be created in the source language, and published to a repository or local folder.

Context

This procedure presents how to translate the data foundation and business layer metadata using local files. You can also translate metadata in shared projects by accessing the metadata files in the shared project in the repository from the translation management tool. For details on the procedures in the translation management tool, see the *Translation Management Tool User Guide*.

Procedure

1. In the information design tool, create a local project if you do not already have one.
When creating the project, note the file path to the directory where the project files are saved in the file system. The default root directory for all projects is workspace.
2. Retrieve the universe into the local project.
The information design tool saves the .dfx and .blx files in the local project. These files correspond to the data foundation and business layer definitions. These are the files that are used as the source for the translations.

Note

For OLAP universes, only a .blx file is saved.
3. In the translation management tool, translate the data foundation metadata (for relational universes):
 - a) Import the .dfx file from project folder in the local file system.
 - b) Translate the metadata.
 - c) Export the translated content to the local file system.For details on these workflows, see the *Translation Management Tool User Guide*.
4. Follow the same procedure as in previous step to translate the .blx file.
5. In the information design tool, to see the translations:
 - a) In the application language preferences, select translated language as the Preferred Viewing Language. Exit and restart the information design tool for the language change to take effect.
 - b) Open the business layer by double-clicking it in the Local Projects View. You can see the translated metadata in the Query Panel. To open the Query Panel, select the **Queries** pane and click **Insert Query**.
6. Re-publish the business layer so that translations are available to universe users.
For universes published to a repository, you can open the Query Panel on the published universe by right-clicking the universe the Repository Resources View and selecting **Run Query**.

Related Information

[Creating a local project](#) [page 79]
[Retrieving a published universe from a repository](#) [page 70]
[Retrieving a published universe from the local file system](#) [page 70]
[Showing values in a data source](#) [page 185]
[Inserting and editing a query in the business layer](#) [page 290]
[Setting languages used by the information design tool](#) [page 32]
[Publishing a universe](#) [page 324]
[Running a query on a universe published in a repository](#) [page 98]

3.9 About the New Universe wizard

You can use the New Universe Wizard to create the resources you need to publish a local universe: either a single-source relational or OLAP universe.

You can also select existing resources. At each step of the process, you can choose to create a resource or select an existing resource.

For more information on each step of the wizard, click the help icon.

To start the wizard, from the information design tool main menu, select **File > New Universe**.

Once you have completed the wizard, publish the universe to a local folder or repository.

Related Information

[Publishing a universe](#) [page 324]
[Publishing a local universe to the repository](#) [page 326]

3.9.1 Select or create a project in the New Universe wizard

Create a project if you want to build all the resources for your universe from scratch. All the resources used to build the universe must be located in the same local project.

If you want to base the universe on existing resources, choose the project that contains these resources.

Related Information

[About local projects and resources](#) [page 78]

3.9.2 Select the data source type in the New Universe wizard

A universe is based on a relational or an OLAP data source.

- Relational universes are based on a data foundation that defines the relevant tables and joins from one or more relational databases. The objects in the business layer map to the database structures via SQL expressions. In the next step, you will select or create a relational connection on which to base the data foundation.
- OLAP universes are based on a connection to an OLAP cube. The objects in the business layer map directly to the cube via MDX expressions. In the next step, you will select or create an OLAP connection on which to base the business layer.

3.9.3 Select or create a relational connection in the New Universe wizard

The wizard lets you to create only local connections. If you want to base your data foundation on a secured connection, do one of the following:

- Select an existing secured connection.
- Cancel the wizard and use the New Relational Connection wizard to create a secured connection.
- Continue with the wizard and create a local connection. You can publish the connection and re-publish the business layer to a repository later.

Related Information

[Creating a relational connection](#) [page 102]

[Publishing a local universe to the repository](#) [page 326]

[About local connections](#) [page 100]

3.9.4 Select or create an OLAP connection in the New Universe wizard

The wizard lets you to create only local connections. If you want to base your universe on a secured connection, do one of the following:

- Select an existing secured connection.
- Cancel the wizard and use the New OLAP Connection wizard to create a secured connection.
- Continue with the wizard and create a local connection. You can publish the connection and re-publish the business layer to a repository later.

Related Information

[Creating an OLAP connection](#) [page 123]

[Publishing a local universe to the repository](#) [page 326]

[About local connections](#) [page 100]

3.9.5 Select or create a data foundation in the New Universe wizard

The wizard lets you to create only single-source data foundations. If you want to create a multisource-enabled data foundation on a secured connection, do one of the following:

- Select an existing multisource-enabled data foundation. In the previous step you must select one of the secured connections referenced in the data foundation.
- Cancel the wizard and use the New Data Foundation wizard to create a multisource-enabled data foundation.

To learn more about single-source and multisource-enabled data foundations, see the related topics.

Related Information

[About single-source data foundations](#) [page 133]

[About multisource-enabled data foundations](#) [page 134]

[How to build a data foundation](#) [page 138]

4 Converting .unv universes

4.1 About .unv and .unx universes

You can use the **Convert .unv Universe** command in the information design tool to convert a universe created with other SAP BusinessObjects universe design tools, and universes created in previous versions. You can then work on the converted universe in a local project as you would a universe created with the information design tool.

What is a .unv universe?

A .unv universe refers to a universe created with any SAP Business Objects XI 3 design tool, for example Universe Designer.

The following SAP Business Objects BI 4 design tools create .unv universes:

- The universe design tool (new name for Universe Designer)
- The universe design tool desktop edition (new name for Universe Designer Personal)

The universe is stored with a file name of `<universe name>.unv` in a local folder, or in a repository.

What is a .unx universe?

When you publish a universe using the information design tool, the universe is stored with a file name of `<universe name>.unx`. This is called a .unx universe. The **Convert .unv Universe** command converts a .unv universe to the .unx universe format.

Related Information

[About converting .unv universes](#) [page 59]

4.2 About converting .unv universes

Before working with .unv universes in the information design tool, you must convert them.

What .unv universes can be converted?

You can convert the following types of .unv universes:

- Relational universes created using the SAP BusinessObjects BI 4 tools universe design tool or universe design tool desktop edition.
- Relational universes created using SAP BusinessObjects Enterprise XI 3 design tools.

Note

Before you can convert universes created in version XI 3 that are saved in a repository, you must upgrade the universes using the upgrade management tool. For more information, see the *SAP BusinessObjects Business Intelligence platform Upgrade Guide*.

You cannot convert the following types of .unv universes:

- OLAP universes
- Stored procedure universes
- Universes based on a Data Federator data source
- Javabeau universes

Note

It is not possible to convert Business Views created with Business View Manager XI 3 to a format compatible with version BI 4 reporting tools.

How to convert .unv universes

How you convert .unv files depends on the software version of the tool that was used to create the universe, and whether the universe is stored locally or in a repository. The following table describes the steps to follow for different conversion scenarios. For details of the conversion procedures, see the Related Topics.

Universe to convert	Workflow
A .unv universe that has been saved in a repository using XI 3 design tools.	<p>First upgrade the universe in the repository to the latest version using the upgrade management tool.</p> <p>In the information design tool, follow the procedure for converting a .unv universe in a repository.</p> <p>The conversion creates an equivalent .unx universe in the repository with the associated universe and connection rights.</p>
A .unv universe that has been saved in a repository using the universe design tool release BI 4.0 or later.	<p>In the information design tool, follow the procedure for converting a .unv universe in a repository.</p> <p>The conversion creates an equivalent .unx universe in the repository with the associated universe and connection rights.</p>

Universe to convert	Workflow
Any locally-stored .unv universe that has been created using version XI 3 or later design tools.	In the information design tool, follow the procedure for converting a locally-stored .unv universe.
<p>i Note</p> <p>A locally-stored universe refers to a non-secured universe that was saved for all users.</p>	<p>The conversion creates the equivalent universe resources (data foundation, business layer, and local connection) in a local project.</p> <p>Publish the business layer to create the .unx universe.</p>

After converting .unv universes

When converting a .unv universe, the .unv universe is preserved. Documents in SAP BusinessObjects query and reporting tools based on the universe are still linked to the .unv universe. This gives you the opportunity to check and test the converted universe before changing the documents that depend on it.

Some features of .unv universes are implemented differently in the .unx universe. Once you have converted a universe, you can edit the universe resources in a local project in the information design tool to check and correct inconsistencies; and to take advantage of new universe features. For a description of the supported features and how they are implemented in .unx universes, see the related topic.

After converting the universe, it is recommended to refresh the structure of the data foundation and run a check integrity on the universe. For tips on resolving check integrity errors on converted universes, see the related topic.

Related Information

[Converting a .unv universe in a repository](#) [page 67]

[Converting a locally-stored .unv universe](#) [page 68]


[Features supported when converting .unv universes](#) [page 61]

[Tips for resolving check integrity errors after converting .unv universes](#) [page 65]

[About .unv and .unx universes](#) [page 59]

4.3 Features supported when converting .unv universes

When converting a .unv universe with the information design tool, the conversion process creates equivalent features in the converted universe. The table below indicates which .unv universe features are supported and how they are implemented in the .unx universe. For certain features, tips are described for obtaining best conversion results.

Feature in original .unv universe	Feature in converted .unx universe
Universe schema	<p>The objects in the universe schema are created in the data foundation:</p> <ul style="list-style-type: none"> • Tables • Alias tables • Derived tables (including nested derived tables) • Joins (including shortcut joins) • Self-joins (converted to column filters) <div> <p>➔ Tip</p> <p>@Prompt functions in self-join expressions may require manual intervention after conversion. See the related topic on resolving check integrity errors.</p> </div> <ul style="list-style-type: none"> • Contexts <div> <p>➔ Tip</p> <p>Contexts are converted with all joins explicitly included or excluded. In the data foundation editor, you can take advantage of the simplified contexts feature. You can manually restrict the context definition to the ambiguous parts of the schema using neutral joins. For more information about contexts, see the related topic.</p> </div> <div> <p> Restriction</p> <p>When converting a .unv universe, SQL is generated in the data foundation for certain object definitions in the .unv universe (for example, table names). If the .unv object definition contains a reference to a business object in an @Prompt, an SQL list of values is generated in the data foundation. The list of values has the following restrictions:</p> <ul style="list-style-type: none"> • If the .unv object is index-aware, the index awareness is not applied. • The Security Access Level of the .unv object is not applied. • Any table mapping restrictions on the .unv object are not applied in the Data Security Profile. </div>
Universe outline	<p>The objects in the universe outline are created in the business layer with all their properties:</p> <ul style="list-style-type: none"> • Classes and sub classes (converted to folders) • Dimensions. For time hierarchies, a dimension is created for each active level in the hierarchy. • measures (including aggregation function) • details (converted to attributes) • conditions (converted to filters, including properties for mandatory filters)
Multilingual universe	All translated strings, language settings, and locale settings are converted.

Feature in original .unv universe	Feature in converted .unx universe
Linked universe (core and derived)	<p>A core universe is converted like any .unv universe. The conversion of a core universe does not trigger the conversion of the derived universes that depend on it.</p> <p>A derived universe contains a link to a core universe. When you convert a derived universe, the conversion automatically includes all core universes that the derived universe is linked to. You do not need to convert the core universes ahead of time as a separate step.</p> <p>The data foundation of the converted universe contains all the tables and joins from all core universes and any tables, joins, or contexts that were defined in the derived universe.</p> <p>The business layer contains all classes, objects, and conditions from all core universes including the classes, objects and conditions that were defined in the derived universe.</p>
Universe Controls Parameters: Query Limits	Query limits are converted and can be edited in the business layer.
Universe SQL Parameters: SQL restrictions	<p>Query, multiple SQL statements, and Cartesian product controls are converted.</p> <p>You can edit the Allow Cartesian products and Multiple SQL statements for each context restrictions in the data foundation. Edit all other SQL restrictions in the business layer.</p>
Strategies	Custom strategies are not supported in .unx universes.
Universe Parameters: SQL generation parameters	<p>Customizations to SQL generation parameter settings in the PRM file or in the universe parameters are not converted. You can add customized values to the converted PRM file, and customize universe parameter settings in the converted universe using the information design tool.</p> <div> <p>➔ Tip</p> <p>Check and reset custom settings to SQL generation parameters in the data foundation properties and business layer properties. For information on setting SQL parameters, see the related topic.</p> </div>
@Functions	<p>The following @functions are converted:</p> <ul style="list-style-type: none"> • @Aggregate_Aware • @Prompt • @DerivedTable • @Select • @Variable • @Where <p>The syntax for all functions is supported.</p>

Feature in original .unv universe	Feature in converted .unx universe
	<p>➔ Tip</p> <p>The @Prompt function has a new alternative syntax to take advantage of named parameters. For more information, see the related topic.</p>
Prompts	<p>@Prompt functions in dimension and measure expressions in the business layer, and in SQL expressions in the data foundation are supported.</p> <p>When converting, you have the option to automatically create a named parameter for the prompt in the business layer.</p> <p>@Prompt expressions in the data foundation are not converted. @Prompt functions in self-join expressions may require manual intervention after conversion. See the related topic on resolving check integrity errors.</p> <p>➔ Tip</p> <p>Parameters and lists of values in the information design tool can be defined independently from the objects that they reference. You can therefore reference a named parameter or list of values in more than one business layer object.</p>
Lists of values	<p>Named lists of values are created in the business layer for dimension and measure objects that specify a list of values.</p>
Access Restrictions	<p>When converting a universe in a repository, universe Access Restrictions are converted into security profiles that can be edited using the Security Editor:</p> <ul style="list-style-type: none"> • Access Restrictions (except object restrictions) are converted to settings in a Data Security Profile. • Object Access Restrictions are converted to Create Query and Display Data settings in a Business Security Profile. <p>➔ Tip</p> <p>Business Security Profiles allow you to secure the metadata separately from the data. For example, you can allow a user to create a query although the user is not allowed to see the corresponding data. For more information about universe security in the information design tool, see the related topic.</p>
Security assignments and priority	<p>When converting a universe in a repository, user and group assignments are converted.</p> <p>➔ Tip</p> <p>In the Security Editor, you can take advantage of the feature allowing you to assign more than one security profile to a user or group.</p>

Feature in original .unv universe	Feature in converted .unx universe
	<p>Group priority for access restrictions is converted.</p> <p>i Note</p> <p>For .unv universes, if a user belongs to different groups, the priority assigned to the groups determines what access restriction the user inherits if he has no access restriction assigned. In the converted universe, priority is assigned to Data Security Profiles instead of groups. If the priority for the profile assigned to the group is higher than the priority of the profile assigned to the user, the group profile is used.</p>
Connections	<p>When converting a universe in a repository, the same secured relational connection is used by both the .unv and .unx universes. If you retrieve the converted universe into a local project, a connection shortcut is created that references the secured connection in the repository.</p> <p>➔ Tip</p> <p>Relational connections can be created and shared by the universe design tool and the information design tool. The connections are published in the same Connections folder in the repository.</p> <p>When converting a locally-stored (non-secured) universe, personal and shared connections are converted to local connections.</p>

Related Information

[Tips for resolving check integrity errors after converting .unv universes](#) [page 65]

[About the Data Foundation Editor](#) [page 136]

[About contexts](#) [page 175]

[About SQL Generation Parameters](#) [page 433]

[About @Functions](#) [page 423]

[About the Business Layer Editor](#) [page 225]

[About universe security](#) [page 328]

4.4 Tips for resolving check integrity errors after converting .unv universes

After converting a .unv universe, it is recommended to run a check integrity in the information design tool on the converted universe. Certain errors in the check integrity results can be resolved by following the best practices described below.

Errors on column data types

Refresh the structure in the data foundation just after the conversion. This avoids data-type errors in the integrity check.

Errors on self-joins with @Prompt

If a join expression in the .unv universe contains an @Prompt with a list of values referring to an object, the converted join needs to be reworked in the data foundation. The steps to follow for two possible solutions are described. The descriptions use the following example:

The .unv universe contains a self-join on the **dimProductStrings** table with a prompt called **Language**. The self-join expression is:

```
dimProductStrings.LanguageID= @Prompt('Language','N','Language\Language Id',mono,constrained)
```

After converting the universe, the data foundation contains a column filter on the **dimProductStrings** table. The join expression for the filter contains the @Prompt.

The first solution consists of creating a prompted parameter and list of values in the data foundation:

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View
2. Click the **Parameters and Lists of Values** tab.
3. In the data foundation **Lists of Values** pane, define a list of values based on custom SQL for Language. For example:

```
SELECT "LANGUAGES"."LANGUAGEID", "LANGUAGES"."LANGUAGECODE" FROM "LANGUAGES"
```
4. In the data foundation **Parameters** pane, define a parameter for Language. Select the option **Prompt to users** and associate the Language list of values to it.
5. In the data foundation, edit the column filter in the **dimProductStrings** table. Change the join expression to reference the new prompted parameter, for example:

```
dimProductStrings.LanguageID = @Prompt(Language)
```
6. Save and close the data foundation.

The second solution consists of using a mandatory filter in the business layer:

1. Open the data foundation in the editor and delete the column filter on the **dimProductStrings** table that contains the @Prompt.
2. Save and close the data foundation.
3. Open the business layer in the editor.
4. In the business layer **Lists of values** pane, define a list of values based on custom SQL for Language. For example:

```
SELECT "LANGUAGES"."LANGUAGEID", "LANGUAGES"."LANGUAGECODE" FROM "LANGUAGES"
```
5. In the business layer **Parameters** pane, define a parameter for **Language**. Keep the default option to **Prompt to users** and associate the **Language** list of values to it.
6. In the business layer, in the folder associated with **Product**, create a filter with an expression that refers to the **Language** prompted parameter, for example:

```
dimProductStrings.LanguageID = @Prompt(Language)
```
7. In the **Properties** tab of the filter definition, select the **Use filter as mandatory in the query** option. Select the **Filter scope** of **Apply on Folder**.

8. Save and close the business layer.

Related Information

[About refreshing a data foundation](#) [page 188]

[About the Data Foundation Editor](#) [page 136]

[About parameters and lists of values in the data foundation](#) [page 180]

[Inserting a column filter](#) [page 164]

[About the Business Layer Editor](#) [page 225]

[About parameters](#) [page 280]

[About lists of values](#) [page 283]

[Inserting and editing filters](#) [page 255]



4.5 Converting a .unv universe in a repository

Prerequisites

The .unv universe to be converted must be stored in a repository compatible with the information design tool. If the .unv universe was created with a design tool version earlier than SAP BusinessObjects BI 4.0, you must first upgrade the universe using the upgrade management tool. For more information about upgrading universes, see the *SAP BusinessObjects Upgrade Guide*.

If you want to retrieve the converted .unx universe into a local project in order to work on it, you must first have a local project folder in the Local Projects View.

Procedure

1. In the information design tool, select **File > Convert .unv Universe** .
2. In the *Convert a .unv Universe* dialog box, click the **Select .unv universe from a repository** icon .
3. Open a session on the repository where the .unv universe is saved, select the universe, and click **OK**.
4. Click the Browse button next to the **Destination Repository Folder** field, and then select a folder in the repository where you want to save the converted .unx universe.
5. If you want to retrieve the converted .unx universe into a local project in order to work on it, click the Browse button next to the **Destination Local Project Folder** field, select a project folder, and click **OK**.
6. If you want the conversion to create named parameters for prompts, select the **Automatically convert @Prompt expressions into universe named parameters** option. For more information about named parameters, see the related topic.

7. If you are retrieving the converted universe into a local project and want remove the local security requirement so that any user can open the universe resources without entering repository authentication, select the **Save for all users** option.
8. Click **OK** to start the conversion.

Next Steps

After conversion, it is recommended to refresh the structure of the data foundation, then run a check integrity on the universe to detect problems in the conversion. For tips on resolving check integrity errors, see the related topic.

Related Information

[Features supported when converting .unv universes](#) [page 61]

[Opening a session](#) [page 96]

[Creating a local project](#) [page 79]

[About parameters](#) [page 280]

[About refreshing a data foundation](#) [page 188]

[Running check integrity](#) [page 316]


[Tips for resolving check integrity errors after converting .unv universes](#) [page 65]

4.6 Converting a locally-stored .unv universe

Prerequisites

You need to define a local project folder in the Local Projects View in which to save the resources of the converted universe.

Procedure

1. In the information design tool, select **File > Convert .unv Universe**.
2. In the *Convert a .unv Universe* dialog box, click the **Select .unv universe from the local file system** icon , and select the universe you want to convert.
3. Click the Browse button next to the **Destination Local Project Folder** field, select a project folder, and click **OK**.

-
4. If you want the conversion to create named parameters for prompts, select the **Automatically convert @Prompt expressions into universe named parameters** option. For more information about named parameters, see the related topic.
 5. Click **OK** to start the conversion.

The conversion creates the equivalent universe resources (data foundation, business layer, and local connection) in the specified local project folder.

At this point, it is recommended to refresh the structure of the data foundation.

You can now publish the business layer to create the .unx universe file. This creates a local universe. To publish the universe to the repository, continue with the next step.
 6. Publish the local connection to a repository.
 7. Edit the data foundation and change the connection to use the secured connection published in the last step.
 8. Publish the business layer to the repository.

Next Steps

The Publishing Wizard allows you to run a check integrity on the universe (recommended). For tips on resolving check integrity errors, see the related topic.

Related Information

[Features supported when converting .unv universes](#) [page 61]

[Creating a local project](#) [page 79]

[About parameters](#) [page 280]

[About refreshing a data foundation](#) [page 188]

[Publishing a local connection to the repository](#) [page 326]

[Changing a connection in a data foundation](#) [page 143]

[Publishing a universe](#) [page 324]

[Running check integrity](#) [page 316]

[Tips for resolving check integrity errors after converting .unv universes](#) [page 65]

5 Retrieving published universes

5.1 Retrieving a published universe from the local file system

Prerequisites

To retrieve a published universe, you must have a project in the Local Projects View where the business layer and referenced resources are to be saved.

Procedure

1. In the Local Projects View, right-click the project folder and select ► **Retrieve Universe** ► **from Local Folder** .
2. Follow the instructions on the wizard pages. For more information about what to do on a particular page, click the help button.

Results

When the wizard finishes, the business layer and its dependent resources (connections, connection shortcuts, and data foundation) are created in the local project and are ready to be edited.

Related Information

[Creating a local project](#) [page 79]

5.2 Retrieving a published universe from a repository

Prerequisites

To retrieve a published universe, you must have a project in the Local Projects View where the business layer and referenced resources are to be saved.

Procedure

1. You can retrieve a universe from a repository in two ways:

Option	Command
From the Local Projects View	Right-click the project folder in the Local Projects View and select Retrieve Universe > from a Repository .
From the Repository Resources View	Right-click the universe in the Repository Resources View and select Retrieve Universe .

Note

By default, the resources are retrieved into the local project and are secured locally by requiring you to enter your repository system authentication when opening a retrieved data foundation or business layer.

To remove the local security requirement, select the **Save for all users** option when selecting the universe in the repository.

2. Follow the instructions on the wizard pages. For more information about what to do on a particular page, click the help button.

Results

When the wizard finishes, the business layer and its dependent resources (connections, connection shortcuts, and data foundation) are created in the local project and are ready to be edited.

Related Information

[Opening a session](#) [page 96]

[Selecting a repository folder](#) [page 325]

[Creating a local project](#) [page 79]

6 Migrating universes to SAP HANA

6.1 About Universe Landscape Migration

Universe Landscape Migration is an add-in to the information design tool that lets you migrate a relational, single-source universe created with the information design tool to a universe that connects to a database on SAP HANA. You can migrate universes based on the following types of relational connections: Oracle, Teradata, Microsoft SQL Server, and Sybase Adaptive Server Enterprise.

The universe's dependent reports (Web Intelligence and Crystal Reports) are also migrated. The security defined on the source universe and reports are applied to the SAP HANA universe and migrated reports.

The Universe Landscape Migration add-in is selected when installing the SAP Business Intelligence Client Tools and the information design tool. For more information, see the *Business Intelligence Platform Installation Guide for Windows*.

For a detailed list of requirements and limitations when using Universe Landscape Migration, see the related topic.

Universe Landscape Migration does the following:

In the pre-migration step:

- Analyzes the source universe and provides a pre-migration report showing universe objects that are impacted during the migration.
- Provides an ATL script file if tables are missing from the SAP HANA database. This script can be used by the Data Services administrator to generate the missing tables.

In the migration step:

- Creates the resources (data foundation and business layer) in a local project for the migrated universe. The data foundation is based on a secured relational connection to SAP HANA that you provide.
- Translates database-specific functions to the equivalent SAP HANA functions.
- Publishes the migrated universe to the repository.
- Migrates the corresponding universe security profiles.

In the post-migration step:

- Migrates the dependent reports that you select and publishes them in the repository.
- Provides a tool to check the differences in results generated by the migrated report versus the original report.

For details on how to migrate a universe to SAP HANA, see the related topics.

Related Information

[Universe Landscape Migration requirements and limitations](#) [page 73]

[Migrating a universe to SAP HANA: Pre-Migration](#) [page 74]

[Migrating a universe to SAP HANA: Migration](#) [page 75]

[Migrating a universe to SAP HANA: Post-Migration](#) [page 76]

6.1.1 Universe Landscape Migration requirements and limitations

Note the following requirements and limitations when migrating universes to SAP HANA:

- The source universe must be single-source and relational. Multisource-enabled and OLAP universes are not supported.
- For proper migration and data comparison, the database schemas, tables, and views used in the source universe should exist in the destination SAP HANA database.
- The user who performs the migration should not have any security profiles assigned to his user name on the universe to be migrated.
- Data foundation tables in the source universe that are based on system tables (for example, DUAL tables in Oracle databases) are not migrated.
- Any SQL function which is not supported by SAP HANA is not migrated.
- Derived tables that use database-specific SQL will not migrate completely. You need to update these tables manually in the migrated universe.
- The schema mapping in the data foundation is done in the order defined in the source database.
- By default user and table names are created in the SAP HANA database in uppercase. If the source database contains two schemas with the same user name, one in uppercase and the other in lowercase, any tables that exist in the lowercase schema are mapped to the uppercase schema in SAP HANA. After migration, tables belonging to the lowercase schema need to be modified in the data foundation. Select these tables and change the schema using the **Change Qualifier/Owner** command.

Note the following requirements and limitations when migrating dependent reports to SAP HANA:

- Only Crystal Reports and Web Intelligence documents are migrated. Dashboards and Explorer documents are not supported.
- When comparing Web Intelligence documents in the report check, if the document has contexts, the check is done by answering the first applicable context for both the documents. This is because Web Intelligence does not remember the previously selected context.
- The report check for Web Intelligence reports having multiple queries is not supported.
- In Web Intelligence documents containing contexts, the first available context is chosen for the migrated report.
- Crystal Reports documents containing contexts are not migrated.
- Crystal Reports documents containing prompts are migrated as blank documents. The user needs to answer the prompt once to provide data for the report.

The following limitations apply to the use of the ATL script file to generate missing tables in the SAP HANA database:

- Use only the following clients for the respective databases:
 - Oracle 10 Oracle Client
 - MS SQL Server 2005 ODBC
 - Teradata 13 ODBC
 - Sybase Adaptive Server Enterprise 15 Sybase Open Client
 - SAP HANA ODBC
- Provide a password in the source and target in Data Services after importing the ATL.
- Tables from multiple databases are not supported.
- Views are created as tables.

6.2 Migrating a universe to SAP HANA: Pre-Migration

Prerequisites



Use this procedure to create the pre-migration reports to help you plan your migration.

Before you begin, you need:

- The universe that you want to migrate. This must be a relational, single-source universe created using the information design tool. The universe must be published in a repository. For a list of the connection types supported for migration, see the related topic about Universe Landscape Migration.
- A secured, relational connection to the SAP HANA server hosting the database. This connection must be published in the same repository as the universe to migrate.

Procedure

1. In the Repository Resources View, open a session on the repository where the universe to be migrated is published.
2. Right-click the universe in the repository, and select **Migrate to SAP HANA**.
If this is the first time you have migrated a universe, the wizard creates a Migration folder in the Local Projects View to contain the migrated resource files.
3. Select a secured connection to SAP HANA:

- Click the  icon to browse the connections in the repository.
- Click  to select a connection shortcut from a local project.

4. Once you have selected a connection, click **Next**.

The wizard creates a connection shortcut in the Migration folder.

The migration preview page shows the following information:

- Any missing items, including missing tables in the SAP HANA database.
 - The objects that will change in the migration, including impacted business layer objects.
 - The list of impacted documents that can be migrated after the universe migration.
5. To generate a report of the pre-migration information, click **Export Report**.
You are prompted for a file path and file name in which to save the PDF report.
 6. If tables are missing in the SAP HANA database, click **Replicate Tables**.
You are prompted for a file path and file name in which to save the ATL report. This report contains an ATL script that the Data Services administrator can use to generate the missing tables in the SAP HANA database.
 7. Review the pre-migration information and decide if want to continue with the migration.
 - If you have errors in the pre-migration report, click **Cancel**. Work with your database administrator to prepare your universe for migration.
 - If you have no errors, you can continue with the migration now. Click **Next**. Details on the migration step of the procedure are in the related topic.

Related Information

[Migrating a universe to SAP HANA: Migration](#) [page 75]

[Opening a session](#) [page 96]

[About Universe Landscape Migration](#) [page 72]

6.3 Migrating a universe to SAP HANA: Migration

Context

Use this procedure to migrate your universe after you have performed the pre-migration step.

If you are continuing from the pre-migration step without cancelling the wizard, start this procedure at step 3.

Procedure

1. Right-click the business layer in the Migration folder in the Local Projects View and select **Migrate to SAP HANA**.
2. Select the connection to SAP HANA: Browse the Migration folder in the Local Projects View, select the connection shortcut and click **Next**.
3. In the SAP HANA migration preview page, select **Next**.

You will be asked to confirm the migration.

The migration summary page shows the following information:

- Objects that changed in the migration.
 - An error log listing errors encountered during the migration.
4. To generate a report of the migration information, click **Export Report**.
 5. To run an integrity check on the migrated resources, click **Check Integrity**.
 6. Review the migration information and decide if you want to publish the migrated universe.
 - If you want to correct errors in the migrated resources, click **Cancel** and edit the migrated resources (the data foundation and business layer) in the Migration folder. Once you have made corrections, restart this procedure.
 - If you want to publish the universe, click **Next**.

The universe is published in the Migrated SAP HANA Universes subfolder in the repository. The subfolder tree of the source universe is duplicated in the Migrated SAP HANA Universes folder.

Once the migrated universe is published in the repository, you can migrate reports now, or exit the wizard and migrate reports later. Details on the post-migration step of the procedure (report migration) are in the related topic.

Related Information

[Migrating a universe to SAP HANA: Post-Migration](#) [page 76]

[Migrating a universe to SAP HANA: Pre-Migration](#) [page 74]

[Running check integrity](#) [page 316]

[About the Data Foundation Editor](#) [page 136]

[About the Business Layer Editor](#) [page 225]

[About Universe Landscape Migration](#) [page 72]

6.4 Migrating a universe to SAP HANA: Post-Migration

Context

Use this procedure to migrate your reports after you have performed the pre-migration and migration steps, and the migrated universe is published to the repository.

If you are continuing from the migration step without cancelling the wizard, start this procedure at step 4.

Procedure

1. Open a session in the Repository Resources View on the repository where the migrated universe is published.
2. In the repository, right-click the migrated universe in the Migrated HANA Universes folder, and select **Post-Migration**.
3. Select the migrated universe and click **Next**.
4. In the Reports page, select the reports you want to migrate.
To generate a report listing all the reports that can be migrated, click **Export Report**.
5. Click **Next** to migrate the selected reports.

The Summary of Migrated reports shows the following information:

- A list of reports that were migrated, including the report type, and the path to where the migrated report is published.
 - An error log listing errors encountered during the migration.
6. To run a comparison of the results of a migrated report with the original report results, select the report and click **Check**.
The status of the comparison appears in the status column of the list of reports. Click the **Error status** to see the details of the errors encountered while comparing the report.
 7. When you are finished checking the migrated reports, click **Finish**.

You can restart the post-migration procedure at any time to migrate other reports.

Related Information

[Migrating a universe to SAP HANA: Migration](#) [page 75]

[About Universe Landscape Migration](#) [page 72]

7 Working with projects

7.1 About local projects and resources

The first step in creating resources in the information design tool is to create a local project in the Local Projects View. You create and edit all resources (except secured connections and security profiles) in a local project.

The resources and folders in a local project are stored as physical files and folders in the local file system. The Local Projects View lets you to navigate local projects and open resources in the information design tool.

Once you have created a local project, there are several ways you can populate it with resources:

- Create universe resources using the wizards available on the **New** menu.
- Convert a .unv universe that was created with the universe design tool, or migrated from an earlier version.
- Retrieve a published universe.
- Create folders to organize resources within the project.
- Create file resources by entering a file name and extension.

You edit the resources using the information design tool editors by double-clicking the resource name in the local project. To open a resource from a list of resources recently opened, select **File > Recent Resources**.

You can get information on the resources you create by right-clicking the resource name and selecting **Properties**. The properties displayed include the path to the resource in the local file system and the date the resource was last modified.

You can also perform the following tasks on resources from the Local Projects View:

- Create a shared project so that you can share resources with other designers.
- Check integrity of data foundations and business layers.
- Edit and test local connections.
- Change and test the connection referenced by a connection shortcut.
- Publish a business layer as a universe to the local file system or a repository.
- Publish a connection to a repository.
- Show dependent resources.
- Save a resource as a report.

When you copy resources, it is best to copy the entire folder so that all the references between the resources are maintained. This is because the paths to referenced resources are relative, not absolute. The information design tool assumes that all resources that reference each other are in the same folder. If you copy a single resource to a location outside of the folder without copying the resources it references, the references are broken.

You can delete a project from the Local Projects View. The project files remain in the local file system until you explicitly delete them. Open the project to make it available again in the Local Projects View.

Related Information


[Creating a local project](#) [page 79]

[Deleting a local project](#) [page 81]

[Opening a local project](#) [page 80]
[About resources in the information design tool](#) [page 19]
[Finding universe resources in the local file system](#) [page 80]
[About converting .unv universes](#) [page 59]
[Retrieving a published universe from a repository](#) [page 70]
[About shared projects](#) [page 82]
[Running check integrity](#) [page 316]
[About publishing resources](#) [page 323]
[About resource dependencies](#) [page 318]
[Saving resources as reports](#) [page 93]

7.1.1 Creating a local project

Procedure

1. From the information design tool main menu, select **File > New > Project**.
2. Give the project a unique name.
3. In **Project Location**, the file path to the default root directory for all projects (workspace) displays. To select a different local folder to contain the project, click the browse button .
4. Click **Finish**.

Results

The project is created in the local file system and displayed in the Local Projects View.

Related Information

[About local projects and resources](#) [page 78]

7.1.2 About resource names

Resource names identify the connections, data foundations, and business layers in the local project. You give a name to the resource when you create it. The name must be unique within the local project.

Note

If you use the same name for resources in different projects, when publishing the resources to the same repository you may have naming conflicts because the names are not unique.

You can optionally enter a description of the resource.

Related Information

[About resources in the information design tool](#) [page 19]

[About resource dependencies](#) [page 318]

7.1.3 Finding universe resources in the local file system

Procedure

1. In the Local Projects View, open the project that contains the universe resources.
2. Right-click a resource (for example, a data foundation, connection, or business layer) and select **Properties**.
The properties displayed include the path to the resource in the local file system and the date the resource was last modified.

Related Information

[About local projects and resources](#) [page 78]

7.1.4 Opening a local project

Context

Information design tool projects saved on the file system can be opened in the Local Projects View.

Procedure

1. From the information design tool main menu, select **File > Open Project**.
2. Select the **Select root directory** option and click **Browse**.
The *Browse for Folder* dialog box opens with the default root directory for all projects (workspace) already selected.
3. Click **OK** to select the default directory, or browse to the folder containing the project you want to open.
All projects not already open in the Local Projects View are listed in **Projects** and are selected by default.

4. Clear the checkbox for any projects you do not want to open and click **Finish**.

Related Information

[About local projects and resources](#) [page 78]

7.1.5 Deleting a local project

Procedure

1. Right-click the project in the Local Projects View and select **Delete**.
2. If you want to permanently delete the project from the Local Projects View and the local file system, select the **Delete project contents on disk** option in the [Confirm Project Delete](#) dialog box.

Note

If you select this option, the deletion is permanent and cannot be undone.

3. Click **Yes** to confirm the deletion.

Next Steps

If you did not delete the project contents permanently, you can open the project to make it available again in the Local Projects View.

Related Information

[Opening a local project](#) [page 80]

7.1.6 Backing-up and recovering universe resources in local projects

Context



Sometimes when the information design tool ends unexpectedly, the local workspace becomes corrupted and you are unable to restart the information design tool. Use this procedure to recreate the workspace and recover your local projects.

Procedure

1. On the local file system, navigate to the folder %USERPROFILE%\businessobjects\bimodeler_14\. For example:
`C:\Documents and Settings\Administrator\businessobjects\bimodeler_14\`
2. Rename the workspace folder to workspace.bak.
3. Start the information design tool.
A new workspace folder is created automatically.
4. Select **File > Open Project** from the main menu.
5. In the *Import Existing Projects* dialog box, select the root directory and browse to the path of the workspace.bak folder created in step 2.
6. Select the projects you want to restore.
7. Select the **Copy projects into workspace** option and click **Finish**.

7.1.7 Searching for and filtering resources in the Local Projects View

Procedure

1. To filter the types of resources displayed in the Local Projects View, click the filter icon  in the icon bar of the view. Select the types to include or exclude.
The Local Projects View displays only resources of the types selected.
2. To search the list, click the **Show/Hide search bar** icon .
3. In the search text box, enter the text and press the `Enter` key to start the search.
The first resource name that contains the search text is highlighted in the view. The total number of resource names containing the search text displays in the search text box.

Note

The search does not highlight resource types that are unselected in the filter.

4. To highlight the next resource found, press the `Enter` key again. Use the `Enter` key to browse through all resource names that match the search text.

7.2 About shared projects

A shared project is a project in a repository whose resources are available to other designers. You create a shared project in a repository from an existing local project in the Local Projects View.

To start working on shared resources, you use the following tasks in the Project Synchronization View:

- Synchronize the project to copy resources between the local and shared projects.
- Lock and unlock resources in the shared project to inform other designers when you are working on them.
- Synchronize a shared project created by another designer. This creates a local project associated with the shared project so that you can start working on the shared resources.

Related Information

[Creating a shared project from a local project](#) [page 83]

[Working in a shared project](#) [page 84]

[Synchronizing a project](#) [page 89]

[Locking a resource](#) [page 90]

[Unlocking a resource](#) [page 91]

[Merging changes to shared resources](#) [page 92]

7.2.1 Creating a shared project from a local project

Prerequisites

The shared project will automatically be given the same name as the local project. A project with this name cannot already exist in the repository. If you need to rename an existing shared project, use the rename command in the Project Synchronization View. To do this, see the related topic.

Procedure

1. In the Local Projects View, right-click the project you want to share and select **New Shared Project**.
2. In the [Open Session](#) dialog box, select the repository system session you want to open and enter your system authentication.
The Project Synchronization View opens showing a shared project with the same name as the local project. At this point, the shared project is empty.
3. In the Project Synchronization View, synchronize the resources you want to save in the shared project.

Related Information

[Renaming a shared project](#) [page 84]

[Opening a session](#) [page 96]

[Synchronizing a project](#) [page 89]

7.2.2 Working in a shared project

Context

Use this procedure to work on resources in an existing shared project.

Procedure

1. Open the Project Synchronization View with a session on the repository system where the shared project is saved.
2. Select the shared project from the **Shared Project** list.
3. In the selected shared project, lock the resources you want to work on.

Locks are available as a communication tool between designers. When other designers open the Project Synchronization View, your lock informs them that you are making changes. It also prevents other designers from updating these resources in the shared project while you have them locked. However, any designer can unlock the resource if necessary.

4. Synchronize the project to update the resources in the local project with the latest changes saved on the server.

If you do not already have a local version of the project, one is created in the Local Projects View.

You may want to review the changes made on the server before updating them in the local project. For more information, see the related topic on merging changes in shared resources.

5. Once you have made your changes, in the Project Synchronization View, synchronize the project to save your changes on the server.
6. Unlock the resources.

Related Information

[Opening the Project Synchronization View](#) [page 88]

[Locking a resource](#) [page 90]

[Synchronizing a project](#) [page 89]

[Merging changes to shared resources](#) [page 92]

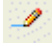
[Unlocking a resource](#) [page 91]

7.2.3 Renaming a shared project

Context

Use this procedure to rename a shared project that exists in the repository.

Procedure

1. Open the Project Synchronization View with a session on the repository system where the shared project is saved.
2. Select the shared project from the **Shared Project** list.
3. Click the **Rename Shared Project** icon .
4. Enter a new name that is unique in the repository.
5. To create a local project with the new name in the Local Projects View, synchronize the project. Local projects with the original name are no longer associated with the shared project with the new name. The resources in those local projects can no longer be synchronized with the newly named shared project.


Related Information

[Opening the Project Synchronization View](#) [page 88]

[Synchronizing a project](#) [page 89]

7.2.4 Deleting a shared project

Procedure

1. Open the Project Synchronization View with a session on the repository system where the shared project is saved.
2. Select the shared project from the **Shared Project** list.
3. Click the **Delete Shared Project** icon .

Note

Deleting a shared project cannot be undone.

Results

The shared project is deleted from the repository. Resources in the local projects associated with the deleted shared project are not affected, however synchronization statuses of the local project are lost.

7.3 About project synchronization

Synchronizing a project starts with comparing the resources in a project in the Local Projects View with an associated shared project on the repository server. Synchronization detects added resources, deleted resources, and differences between the resources. Based on the differences detected, you can update the local and shared resources.




Use the Project Synchronization View to synchronize a project. The view displays synchronization information in two panes:

- The **Shared Project** pane lists the resources in the shared project on the server. A lock icon appears next to the resource if it is locked. Other information about the resources on the server is given: the user who last modified the resource and on what date, the user who locked the resource and on what date.
- The **Synchronization Status** pane lists the status of each resource. The status is determined by comparing the resources in the local and shared projects.

The different synchronization statuses and what they mean are listed in the table.






Status	Description
Added Locally	The resource was added in the local project, but not in the shared project.
Changed Locally	The resource was changed in the local project, but not in the shared project since the last synchronization.
Deleted Locally	The resource was deleted in the local project but still exists in the shared project.
Added on Server	The resource is not in the local project but exists in the shared project.
Changed on Server	The resource was changed in the shared project, but not in the local project since the last synchronization.
Deleted on Server	The resource exists in the local project, but was deleted in the shared project.
Conflicting	Any of the following situations creates a conflicting status: <ul style="list-style-type: none">• The resource was changed in both the local and shared projects with different changes since the last synchronization.• A resource with the same name was added both in the local and shared projects since the last synchronization.• The resource was changed in the local project, but deleted from the shared project.• The resource was changed in the shared project, but deleted from the local project.
Synchronized	The resources are identical.

Three commands allow you to synchronize resources. When you select the resources to be synchronized, you can select individual resources or folders. The following table summarizes the possible synchronization actions.

Icon	Command	Synchronization Action
	Get Changes from Server	<p>For the selected resources:</p> <ul style="list-style-type: none"> If the status is Added on Server, the resource is added to the local project. If the status is Changed on Server, the resource is updated in the local project. If the status is Deleted on Server, the resource is deleted from the local project. If the status is Conflicting, the resource on the server (whether it is changed, added, or deleted) is copied to the local project, regardless of the change made in the local project. <p>For all other statuses, no action is taken.</p> <div> <p>i Note</p> <p>You may want to review the changes made on the server before updating them in the local project. For more information, see the related topic on merging changes in shared resources.</p> </div>
	Save Changes on Server	<p>For the selected resources:</p> <ul style="list-style-type: none"> If the status is Added Locally, the resource is added to the shared project on the server. If the status is Changed Locally, the resource is updated in the shared project on the server. If the status is Deleted Locally, the resource is deleted from the shared project on the server. If the status is Conflicting, the resource in the local project (whether it is changed, added, or deleted) is copied to the shared project, regardless of the change made in the shared project. <div> <p>i Note</p> <p>If a resource is locked by another user, an error message displays and the changes and deletions are not made on the server.</p> </div> <p>For all other statuses, no action is taken.</p>
	Revert Changes	<p>For the selected resources, the local project is updated with the shared project on the server, regardless of the status.</p>

Icon	Command	Synchronization Action
		<p>i Note</p> <p>Revert Changes updates the local project in the same way as Get Changes from Server with the exception that if a resource has been created in the local project and has not yet been saved on the server, Revert Changes deletes the new local resource, whereas Get Changes from Server preserves the new local resource.</p>

The list of resources in the **Synchronization Status** pane can be filtered on status using the icons in the pane's tool bar:

	Shows all resources. This clears the filters and lists all resources regardless of their status.
	Shows/Hides the resources with a status of Synchronized .
	Shows/Hides the resources in the local project that have changed with respect to the server.
	Shows/Hides the resources with a status of Conflicting .
	Shows/Hides the resources on the server that have changed with respect to the local project.

Related Information

[Synchronizing a project](#) [page 89]


[Locking a resource](#) [page 90]

[Unlocking a resource](#) [page 91]

[Merging changes to shared resources](#) [page 92]

7.3.1 Opening the Project Synchronization View

Procedure

1. To open the Project Synchronization View, from the main menu, select **Window > Project Synchronization**.
2. Click the **Change Session**  icon in the Project Synchronization View to open a session on the repository system where the shared projects are saved.

Results

After entering your authentication information, you can manage shared projects, and select a project to synchronize in the **Shared Project** list.

Related Information

[Opening a session](#) [page 96]

[About project synchronization](#) [page 86]

7.3.2 Synchronizing a project

Prerequisites



To synchronize a project, the project must be shared.

Context

Synchronize a project when you want do the following tasks:

- Update local resources with changes stored in the shared project.
- Save in the shared project the changes you made to local resources.
- Revert local resources to the copy stored in the shared project.
- Create a local copy of a shared project.


Procedure

1. Open the Project Synchronization View with a session on the repository system where the shared project is saved.
2. Select the shared project from the **Shared Project** list.
3. To see the latest synchronization status of the resources in the project, in the pane labeled **Synchronization status (Local Project Compared to Shared Project)**, expand the project and click the refresh icon .
For more information about synchronization status and possible actions, see the related topic about project synchronization.
4. Synchronize the project:
 - To update the local project with resources that have changed in the shared project, select the resources in the list and click the **Get Changes from Server** icon .

If no project exists in the Local Projects View with the name of the shared project, a local project is created.

Note


You may want to review the changes made on the server before updating them in the local project. For more information, see the related topic on merging changes in shared resources.

- To update the shared project with changes made locally, select the resources in the list and click the **Save Changes on Server** icon. .

Note

Resources that are locked by another user cannot be updated on the server. However, if needed, any user can unlock the resource.

If you are updating the server with resources that you have locked, synchronizing updates the resources, but does not unlock them. You must explicitly unlock the resources on the server.

- To revert resources in the local project to the copy stored on the server, select the resources and click the **Revert Changes** icon. .

Related Information

[About project synchronization](#) [page 86]

[Opening the Project Synchronization View](#) [page 88]

[Locking a resource](#) [page 90]

[Unlocking a resource](#) [page 91]

[Merging changes to shared resources](#) [page 92]

7.3.3 Locking a resource

Prerequisites

To lock a resource, the resource must be in a shared project.

Context

Lock a resource when you want to inform other designers that you are working on the resource when they open the Project Synchronization View.

Procedure

1. Open the Project Synchronization View with a session on the repository system where the shared project is saved.
2. Select the shared project from the **Shared Project** list.
3. In the **Shared Project** pane, expand the project.
4. Right-click the resource and select **Lock**.

Note

The lock action does not update the resource contents in either the local or shared project. To save any changes, synchronize the resource.

Related Information

[Opening the Project Synchronization View](#) [page 88]

[About project synchronization](#) [page 86]

7.3.4 Unlocking a resource

Context

Unlock a resource once you have updated your changes on the server and you want to inform other designers that you have finished. Once you have unlocked the resource, other designers can lock it and/or update the server version with changes.

Note

If necessary, you can unlock a resource locked by another user.

Procedure

1. Open the Project Synchronization View with a session on the repository system where the shared project is saved.
2. Select the shared project from the **Shared Project** list.
3. In the **Shared Project** pane, expand the project.
4. Right-click the resource and select **Unlock**.

Note

The unlock action does not update the resource on the server with any changes made in the local project. To save any changes, synchronize the resource.

Related Information

[Opening the Project Synchronization View](#) [page 88]

[About project synchronization](#) [page 86]

7.3.5 Merging changes to shared resources

Context

When synchronizing a shared resource, before getting changes from the server, you may want to review the changes and decide which changes to apply to the local resource. This procedure presents a way to manually merge changes between differing resources.

For example, you are working on a resource called **NewDatafoundation** in a local project. This project is shared in a repository. When you synchronize the project, **NewDatafoundation** has a synchronization status of **Changed on Server**, or **Conflicting**.

To review and manually merge the changes:

Procedure

1. In the local project, right-click **NewDatafoundation** and select **Copy**.
2. Right-click again (in the local project) and select **Paste**.
A copy of **NewDatafoundation** is saved in the local project.
3. In the Project Synchronization View, select **NewDatafoundation** and synchronize by selecting **Get Changes from Server**.
4. Open both **NewDatafoundation** and **Copy of NewDatafoundation** in the Data Foundation Editor by double-clicking each resource name in the local project.
Each copy will open in a separate tab of the editor.
5. Compare the changes from the server in **NewDatafoundation** to your local changes in **Copy of NewDatafoundation**.
6. In the editor tab with **NewDatafoundation** open, delete any server changes you do not want to keep, and add any changes that you made locally that you want to keep.
7. Save the changes to **NewDatafoundation** in the editor.

-
8. In the Project Synchronization View, refresh the synchronization. Update the server by selecting **NewDatafoundation** and then selecting **Save Changes on Server**.

Next Steps

Eventually, once you have verified the merged changes, you can delete **Copy of NewDatafoundation** from the local project.


7.4 Saving resources as reports

Context

You can save any resource in a local project as a report in a local file.

Procedure

1. Right-click the resource name in the Local Projects View, and select **Save As**.
2. In the **Report Location** box, enter a file path, file name, and file type for the report. The file type can be .pdf, .html, or .txt.

To browse the local file system to find a file path, click the browse button .

3. For larger resources (data foundations and business layers), you can select which metadata elements to include in the report in the **Metadata Elements** box.
4. Click **Generate** to create the report.

8 Working with repository resources

8.1 About managing repository resources

Repository resources are the universes and connections that have been secured in a repository on a Central Management System (CMS). The Repository Resources View lets you navigate and interact with the folders and resources in repositories.

The Connections folder contains the secured connections created using the information design tool and universe design tool.

Note

A CommonConnections sub-folder sometimes appears in the Connections folder. The CommonConnections folder contains OLAP connections created in the Central Management Console for use in SAP BusinessObjects Advanced Analysis.

The Universe folder contains universes published with the information design tool (.unx universes), as well as universes created and exported with the universe design tool or migrated from earlier versions (.unv universes).

To navigate a repository, open a session on the CMS where the repository is stored. For more information about sessions, see the related topic.

The following sections summarize the tasks you can do from the Repository Resources View.

Folder management

With appropriate rights, you can insert, rename, and delete sub-folders in the Connections and Universes folders.

Secured Connection management

- Edit an existing connection.
- Insert a new secured relational or OLAP connection in the repository.
- Create a connection shortcut in a local project from an existing secured connection.
- Delete a secured connection from the repository.

Universe management

You can do the following tasks on .unx universes (published using the information design tool):

- Run a check integrity.

- Run a query. This command opens the Query Panel. The security settings defined in the security profiles for the universe are applied according to the username in the session.
- Retrieve a universe. This command saves the business layer and its referenced resources in a local project so that you can edit them.
- Rename a universe. This command renames only the universe, and not the underlying business layer.
- Delete a universe from the repository.

You can do the following tasks on .unv universes (created using the universe design tool or migrated from earlier versions):

- Convert a universe. You can save the converted resources in a local project, or publish the converted .unx universe in the repository.
- Delete a universe from the repository.

Related Information

[About session management](#) [page 95]

[About the Connection Editor](#) [page 102]

[Creating a relational connection](#) [page 102]

[Creating an OLAP connection](#) [page 123]

[About connection shortcuts](#) [page 101]

[Running check integrity](#) [page 316]

[Running a query on a universe published in a repository](#) [page 98]

[Retrieving a published universe from a repository](#) [page 70]

[About converting .unv universes](#) [page 59]

8.2 About session management

A session contains the Central Management Server (CMS) system name and authentication information needed to access resources stored in a repository. You need at least one session defined to connect to a repository. You can define additional sessions that connect to the same repository as a different user.

Workflows in the information design tool that require access to secured resources prompt you with the [Open Session](#) dialog box. If you have not already defined a session for the repository you want to access, you can select **New Session** from the **Sessions** list. You can also define a session with the **Insert Session** command in the Repository Resources View.

Once the session is defined, it is retained in the Repository Resources View, and also on the **Sessions** list. The next time you open the session, you will only need to enter the password.

Once a session is opened it stays open until you exit the information design tool. To explicitly close a session, you must do so from the Repository Resources View.

Several sessions can be open at the same time, as long as the sessions are on different CMS systems. If you need to open a session with a different username and password on a CMS that has another session open, you must first close the open session.

If you no longer need a session and want to delete it from the list, use the **Delete Session** command in the Repository Resources View.

Related Information

[CMC rights for information design tool users](#) [page 330]

[Opening a session](#) [page 96]

[Closing a session](#) [page 97]

8.2.1 Opening a session

Context

Different workflows require you to open a session. If you are prompted to open a session, a list of predefined sessions is available. The **Sessions** list is organized in the following order:

- Open sessions in alphabetical order
- Closed sessions in alphabetical order
- **New Session**

To open a session on a repository already defined in the information design tool:


1. Do one of the following:
 - In the Repository Resources View, right-click the repository name and select **Open Session**.
 - Select the session in the **Sessions** list.
2. The authentication information for the CMS is filled in for you. If the session is not already open, enter your **Password**.

Note

If you try to open a session on a repository which has another session already open, an error message displays. To change sessions on a repository, you must first close the open session in the Repository Resources View.

3. Depending on the workflow, click **OK**, **Next**, or **Connect**.

To open a session on a repository not already defined in the information design tool:

1. Do one of the following:
 - In the Repository Resources View, from the **Insert**  menu, select **Insert Session**.
 - Select **New Session** from the **Sessions** list.
2. In the **System** box, enter the name of the Central Management System where the repository is located.

Note

To insert a session for a repository that is hosted on a machine in a different domain than the client hosting the application, you need to provide the host information in a hosts file on the client. Update the hosts file in the following location:

```
C:\WINDOWS\system32\drivers\etc\hosts
```

3. Enter your **User Name** and **Password**.

Note

For **Authentication** type **Windows AD**, specify the complete domain name in the **User Name**. For example, enter **myuser@domain.com** instead of **myuser@domain**.

4. In the **Authentication** list, select the authentication method to use.
5. Depending on the workflow, click **OK**, **Next**, or **Connect**.

Results

The session remains open until you explicitly close it in the Repository Resources View, or exit the information design tool.

Related Information

[Closing a session](#) [page 97]

8.2.2 Closing a session

Context

All open sessions close when you exit the information design tool. To explicitly close a session:

Procedure

1. In the Repository Resources View, select the session you want to close.
2. Right-click and select **Close Session**.

8.3 Running a query on a universe published in a repository

Context

When you run a query on a universe published in a repository, the Query Panel applies the settings defined in the security profiles for the universe according to the user name defined in the session.

Procedure

1. In the Repository Resources View, select the universe. Select .unx universes only.
2. Right-click the universe name and select **Run Query**.

Results

The Query Panel opens with a list of views and objects granted for your user name.

Related Information

[How to build a query](#) [page 297]

9 Working with connections

9.1 About connections

A connection is a named set of parameters that define how one or more SAP BusinessObjects applications can access relational or OLAP data sources. The connection can be a local file, or a remote object in a repository that is referenced by a local shortcut in the information design tool.

You use connections for the following purposes:

Purpose	Description
Relational data source for a data foundation	You associate one or more relational connections to a data foundation, and build a business layer on the data foundation. When you publish the business layer as a universe, the connections and data foundation are integrated in the universe and provide the data for queries run against the universe.
OLAP data source for a business layer	For an OLAP data source, you associate a business layer directly to a connection. The business layer is published as a universe, but the connection provides direct access to the cube.
Direct access to an SAP NetWeaver BW BEx Query	You define SAP NetWeaver BW connections that use the SAP BICS Client middleware driver to provide access to a BEx Query. SAP BusinessObjects query and reporting applications connect directly to the BEx Query. You cannot use these connections as a source for business layers or universes. For information on how to build a universe on SAP NetWeaver BW, see the related topic on using SAP NetWeaver BW data sources.
Direct access to an SAP HANA information model	You define SAP HANA connections that use the SAP HANA Client middleware driver to provide direct access to a single information model (such as an Analytic View or Calculation View). SAP BusinessObjects query and reporting applications connect directly to the cube representing the information model. You cannot use these connections as a source for business layers or universes. For information on how to build a universe on SAP HANA, see the related topic on using SAP HANA data sources.

Connections can be either local or secured.

Related Information

[About local connections](#) [page 100]

[About secured connections](#) [page 100]

[Using SAP NetWeaver BW data sources](#) [page 40]

[Using SAP HANA data sources](#) [page 45]
[Creating a relational connection](#) [page 102]
[Creating an OLAP connection](#) [page 123]
[About the Connection Editor](#) [page 102]

9.1.1 About local connections

You create local connections in the information design tool local project. Local connections are saved as independent objects on the local file system as .cnx files.

Local connections are used for the following purposes:

- To access relational data sources when authoring a data foundation and relational business layer.

Note

To create a multisource-enabled data foundation, you must reference secured connections.

- To access an OLAP cube when authoring an OLAP business layer.
- To run queries on a target database to test modifications in the business layer, or to build lists of values.

Local connections have limited or no security as they can be used by any user with access to the machine running the information design tool.

To secure a local connection, you publish the connection to a repository.

Related Information

[Publishing a local connection to the repository](#) [page 326]
[About secured connections](#) [page 100]
[Creating a relational connection](#) [page 102]
[Creating an OLAP connection](#) [page 123]
[Editing local and secured connections](#) [page 130]

9.1.2 About secured connections

A secured connection is a connection that has been created in, or published to, a repository. It is stored in a dedicated Connections folder in the repository. You can create subfolders in the Connections folder to organize the storage of connections in the repository.

When a connection is published, a connection object containing the same parameters as the local connection is created in the Connections folder or subfolder in the repository.

You can also create secured connections directly in the repository using the **Insert Relational Connection** and **Insert OLAP Connection** commands from the Repository Resources View.

Secured connections cannot be copied to the local file system, but are made available in the Local Projects View as connection shortcuts. The shortcut can be used in the same way as a local connection. However, the connection properties can only be modified by connecting to the repository system.

Secured connections and connection shortcuts are used for the following purposes:

- To retrieve data for universes published to a repository.
- To retrieve data for SAP BusinessObjects reporting products directly accessing database middleware.
- As a data source when creating a data foundation or OLAP business layer.

A secured connection is subject to the following general security constraints in the repository:

- Users must be authenticated.
- User rights can be defined at the user level to grant or deny access to connections or connection properties.
- Connections can only be shared and used by authenticated users.

Downloading relational connections locally

To maintain confidentiality, some sensitive secured connection parameters, for example username and password, remain stored in the repository.

In order to edit the connection in the information design tool, the [Download connection locally](#) right must be granted in the Central Management Console (in addition to the [Create, modify, or delete connections](#) application right and the [Edit objects](#) connection right).

If the [Download connection locally](#) right is granted, you can choose to run queries on the server using the server middleware driver, or locally using the local middleware driver. To use the local middleware, select the local middleware option in the information design tool preferences. If this right is denied, the information design tool uses the server middleware.

Related Information

[About connection shortcuts](#) [page 101]

[Publishing a local connection to the repository](#) [page 326]

[CMC rights for information design tool users](#) [page 330]

[Setting middleware for secured relational connections](#) [page 33]

[Editing local and secured connections](#) [page 130]

9.1.3 About connection shortcuts

A connection shortcut is an object that references a secured connection in a repository. The shortcut is saved as a .cns file on the local file system. The shortcut contains the repository address and port number, the type of connection (OLAP or relational), and an ID that identifies the connection on the server.

You use a connection shortcut when authoring, or modifying any data foundation or business layer that uses a connection stored in the repository.

You can create a connection shortcut in two ways:

- Publish a local connection to the repository.
- Create a shortcut from an existing secured connection in the Repository Resources View.

Related Information

[Publishing a local connection to the repository](#) [page 326]

[Creating a connection shortcut](#) [page 129]

[Editing connection shortcuts](#) [page 131]

9.2 About the Connection Editor

Use the Connection Editor to do the following tasks. For more information, see the related topics.

- Edit connection properties and parameters, and change the middleware driver.
- Edit the properties of connection shortcuts, and change the referenced connection.
- Browse the values in the tables referenced by a relational connection.
- Browse the objects in an OLAP cube and run an MDX query on the cube.

Related Information

[Editing local and secured connections](#) [page 130]

[Editing connection shortcuts](#) [page 131]

[Showing values in a relational connection](#) [page 131]

[Showing values in an OLAP connection](#) [page 132]

9.3 Creating a relational connection

Prerequisites

Use the New Relational Connection wizard to create local and secured connections to a relational data source.

Make sure the middleware driver is configured for the data source to which you want to create a connection. For more information about middleware configuration, see the *Data Access Guide*. For information about supported data sources, see the SAP Business Objects BI Platform 4.1 Supported Platforms (PAM) at [http://](http://service.sap.com/pam)

service.sap.com/pam .

Before you create a local connection, you must have a project available in the Local Projects View. For more information on creating local projects, see the related topic.

i Note

You must create relational connections to SAP NetWeaver BW and SAS sources directly in the repository as secured connections.

Procedure

1. Do one of the following:
 - To create a local connection, select the project folder in the Local Projects View. Select **File** **New** **Relational Connection**.
 - To create a secured connection, in the Repository Resources View, open a session on the repository where you want to create the secured connection. Right-click the Connections folder or subfolder in the repository, and select **Insert Relational Connection**.
2. Follow the steps in the New Relational Connection wizard to enter the following information:
 - Name for the connection
 - Middleware for the target database
 - Login parameters to connect to the relational data source
 - Configuration and custom parameters to optimize the connection

If you need help on a particular step, click the help icon in the wizard dialog box.

Related Information

[Creating a local project](#) [page 79]

[About managing repository resources](#) [page 94]

[Name a connection](#) [page 103]

[About connections](#) [page 99]

9.3.1 Name a connection

This section describes the Resource name page of the New Relational Connection wizard.

You name a connection and can enter a description of the data source. The name and description are available as properties of the connection and can be edited at any time.

Properties	Description
Resource name	Connection name. This field is mandatory.

Properties	Description
Description	Information describing data source. This information can be useful when the connection is used for multiple data foundations. This is optional information.

When you have entered name information, click **Next** to continue the wizard.

Related Information

[Select a middleware driver](#) [page 104]

9.3.2 Select a middleware driver

You select a connection driver to connect to the correct middleware version for the target database. The connection driver is the SAP BusinessObjects driver that maps information in the middleware to the user interface of the SAP BusinessObjects application.

Expand the database and middleware node for the target database and select the connection driver. Click **Next** to continue the wizard.

Note

SAP NetWeaver BW and SAS connections are only listed if you are creating the connection directly in the repository.

Note

If you are using the information design tool from a Crystal Server 2011 installation, SAP middleware drivers are not available.

Related Information

[Set the connection parameters](#) [page 104]

9.3.3 Set the connection parameters

The connection parameters vary depending on the type of data source for which you are defining the connection. Select from the related topics the link to more information about the connection parameters.

Related Information

[Login parameters for relational connections](#) [page 105]

[Login parameters for SAP NetWeaver BW and ERP connections](#) [page 107]

[Login parameters for SAS connections](#) [page 110]

[Login parameters for Oracle EBS connections](#) [page 112]

[Login and schema parameters for CSV file connections](#) [page 118]

[Parameters for OData connections](#) [page 113]

[Parameters for XML and Web Services connections](#) [page 114]

[Configuration parameters for relational connections](#) [page 116]

[Custom parameters for relational connections](#) [page 118]

9.3.3.1 Login parameters for relational connections

The following login parameters apply to most relational connections.

Follow the links for a description of login parameters for the following types of connections:

- [Login parameters for SAP NetWeaver BW and ERP connections](#) [page 107]
- [Login parameters for SAS connections](#) [page 110]
- [Login parameters for Oracle EBS connections](#) [page 112]
- [Parameters for OData connections](#) [page 113]
- [Parameters for XML and Web Services connections](#) [page 114]

Parameter	Description
Authentication Mode	<p>The method used to authenticate the user's login credentials when accessing the data source:</p> <ul style="list-style-type: none">• Use specified user name and password: Uses the User Name and Password parameters defined for the connection.• Use BusinessObjects credential mapping: Uses the database credentials associated with the user's account defined on the Central Management Server (CMS) to connect to the data source. The database credentials are set in the User's Properties in the Central Management Console. For more information, see the <i>SAP BusinessObjects Business Intelligence platform Administrator's Guide</i>.• Use Single Sign On: This authentication mode is used to support end-to-end single sign on defined in the Central Management Server (CMS). If you use an external authentication source (for example LDAP), the CMS and the data source must be configured to use this external authentication source. For more information on single sign on, see the <i>SAP BusinessObjects Business Intelligence platform Administrator Guide</i>.
User Name	<p>The user name to access the data source if Authentication Mode is Use specified user name and password.</p>

Parameter	Description
Password	The password to access the data source if Authentication Mode is Use specified user name and password .
Use SSL	This parameter applies only to SAP HANA connections. If selected, uses SSL protocol to connect to the server.
Single Server	This parameter applies only to SAP HANA connections. Check this option if you are connecting to only one SAP HANA database server. <ul style="list-style-type: none"> • Host Name: the name of the server hosting the data source. Do not include the port number. • Instance Number: the SAP HANA instance number, which represents the second and third digits of the port number. Must be set between 00 and 99. For example, if the port number is 30215, the instance number is 02.
Multiple Servers	This parameter applies only to SAP HANA connections. Check this option to take advantage of the SAP HANA failover mechanism. Server (host:port{;host:port}): a list of servers separated by a semi-colon, for example: (host1:30015;host2:30015;host3:30015) . The JDBC driver will choose one of these hosts for connecting. If a host is not available, the driver chooses the next host from the list. You can also enter the host and port of only one server in the field.
Server (<host>:<port>)	The name and port of the server hosting the data source. For connections to Oracle, you can enter a list of servers separated by commas: (<host>:<port>, <host>:<port>).
Server	The name of the server hosting the data source.
Database	The database name.
Data Source Name	For ODBC connections, the name of the data source that you have defined using your operating system's data source manager.
Alias	For DB2 connections, this is the alias of the database that you have created in the DB2 configuration assistant.
Net Service	For Oracle connections using JDBC middleware, the Oracle Net Service name.
Service	For Oracle connections, the alias that contains the server IP and Net Service information.

Parameter	Description
JDBC_URL JDBC_CLASS	For generic JDBC connections, the JDBC URL and class used to connect to the database.
Informix Server	For Informix connections, the name of the Informix server that you have defined.
OLE DB Provider Name	For generic OLE DB providers, the provider name.

9.3.3.2 Login parameters for SAP NetWeaver BW and ERP connections

The following parameters apply to connections to SAP NetWeaver BW (relational and BICS Client connections), and SAP ERP.

To set ABAP function and InfoSet parameters for SAP ERP connections, after entering the login parameters, click **Next**.

Parameter	Description
Authentication Mode	<p>The method used to authenticate the user's login credentials when accessing the data source:</p> <ul style="list-style-type: none"> • Use specified user name and password: Uses the User Name and Password parameters defined for the connection. • Use BusinessObjects credential mapping: Uses the database credentials associated with the user's account defined on the Central Management Server (CMS) to connect to the data source. The database credentials are set in the User's Properties in the Central Management Console. For more information, see the <i>SAP BusinessObjects Business Intelligence platform Administrator's Guide</i>. • Use Single Sign On: This authentication mode is used to support end-to-end single sign on defined in the Central Management Server (CMS). If you use an external authentication source (for example LDAP), the CMS and the data source must be configured to use this external authentication source. For more information on single sign on, see the <i>SAP BusinessObjects Business Intelligence platform Administrator Guide</i>.
Client Number	The number used to identify the client on the SAP system.
User Name	The user name to access the data source if Authentication Mode is Use specified user name and password .
Password	The password to access the data source if Authentication Mode is Use specified user name and password .

Parameter	Description
Language	<p>The 2-character ISO language code for the language to be used for the connection to the data source. For example, EN for English.</p> <div> <p>i Note</p> <p>In some cases, select the language from the list.</p> </div>
Save Language	<p>Specifies which language to use for the connection:</p> <ul style="list-style-type: none"> If you select the Save Language option, the value from the Language parameter is used. If you clear Save Language, the value from the user's session (Preferred Viewing Locale) is used.
System ID	<p>The 3-character SAP system ID.</p> <div> <p>i Note</p> <p>Required for both application and message server types.</p> </div> <div> <p>i Note</p> <p>For a successful connection to the message server, you need to add the message server system ID to the following file on the machine hosting the application:</p> <p>C:\WINDOWS\system32\drivers\etc\services</p> <p>At the end of the existing file, add the line:</p> <p>sapmsXXX <tab> 3601/tcp</p> <p>where sapms means SAP message server, xxx is the is the System ID of the server that is used, and 3601/tcp is the default TCP port used for communication.</p> </div>
Server type	<ul style="list-style-type: none"> Select Application Server to connect directly to the SAP server without using load balancing. Select Message Server to benefit from SAP load balancing capabilities.
Server Name for Application Server	The name of the SAP application server.
System Number for Application Server	The system number of the SAP application server. This is a two-digit integer between 00 and 99.
Server Name for Message Server	The name or IP address of the SAP message server used for load balancing.

Parameter	Description
Group Name for Message Server	Name of the Logon group; a set of dedicated application servers used to log on.

The following parameters apply only to SAP NetWeaver BW connections:

Parameter	Description
Use Custom ID Program Mapping	<p>An optional parameter for SAP NetWeaver BW relational connections only.</p> <p>The Program ID Mapping defines the program IDs for the callback that SAP NetWeaver BW uses to contact the data federation server. Enter Program ID Mapping as a list of one or more server name=program ID pairs separated by the semi-colon character (;). For example:</p> <p><code><MySIA.DF_Server1>=RFC1;<MySIA.DF_Server2>=RFC2</code></p> <p>Each program ID must match the name of an RFC destination created on SAP NetWeaver BW.</p> <p>If this parameter is not defined, the data federation server automatically creates an RFC destination.</p> <p>For more detailed information, see the description of the programIDMapping connector property in the Data Federation Administration Tool Guide.</p>
Use Custom Gateway	<p>An optional parameter for SAP NetWeaver BW relational connections only.</p> <p>In Gateway Host Name, enter the name of the server hosting the SAP NetWeaver BW gateway.</p> <p>In Gateway Service Name, enter the name or port number of the SAP NetWeaver BW gateway service.</p> <p>If this option is not selected, SAP NetWeaver BW provides the gateway host name and service name via an RCF.</p>
InfoProvider	For SAP NetWeaver BW relational connections, the name of the InfoCube or MultiProvider to be used as the fact table in the center of the snowflake schema in the data foundation.
Catalog	<p>For SAP NetWeaver BW relational connections, the name used to identify the connection to the query server.</p> <div> <p>i Note</p> <p>A default catalog name is registered automatically with the query server the first time the connection is added to any multisource-enabled data foundation.</p> </div>

9.3.3.2.1 InfoProvider fact table selection

When you create an SAP NetWeaver BW connection, the [Select InfoProvider Fact Table](#) dialog box lets you choose a fact table that will become the center of the snowflake schema in your data foundation.

The [filter](#) button lets you filter by types of InfoProvider.

9.3.3.2.2 ABAP function and InfoSet parameters for ERP connections


The following parameters apply to SAP ERP connections. For more information about SAP ERP connections, see the *Data Access Guide*.

Parameter	Description
Function Name Wildcard	<p>Wildcards are filters that reduce the number of tables exposed in the connection. The wildcard character is *, and it matches 0 to any number of characters. The wildcard character can be used with keywords. For example:</p> <p>*keyword_one*keyword_two*</p> <p>The above wildcard will expose only tables containing keyword_one followed by keyword_two.</p>
Map Table Parameters into Input Columns	<p>If selected, the table parameters are considered as both input and output parameters of the ABAP function.</p> <p>If unselected, the table parameters are considered as output only.</p>
Map Selection Fields into Table Columns	<p>If selected, any selection field in the SAP Query is mapped into a table column and considered as an optional input column:</p> <ul style="list-style-type: none">• A query can contain only the EQUAL filter in this column• If the column is only in the projection, NULL is returned <p>If unselected, the selection fields are ignored. No filtering is possible on these fields.</p>

9.3.3.3 Login parameters for SAS connections

The following parameters apply to connections to SAS data sources.

To include access to multiple data sets that are not pre-defined to the SAS/SHARE server, after entering the login parameters, click **Next**.

Parameter	Description
Authentication Mode	<p>The method used to authenticate the user's login credentials when accessing the data source:</p> <ul style="list-style-type: none"> • Use specified user name and password: Uses the User Name and Password parameters defined for the connection. • Use BusinessObjects credential mapping: Uses the database credentials associated with the user's account defined on the Central Management Server (CMS) to connect to the data source. The database credentials are set in the User's Properties in the Central Management Console. For more information, see the <i>SAP BusinessObjects Business Intelligence platform Administrator's Guide</i>. • Use Single Sign On: This authentication mode is used to support end-to-end single sign on defined in the Central Management Server (CMS). If you use an external authentication source (for example LDAP), the CMS and the data source must be configured to use this external authentication source. For more information on single sign on, see the <i>SAP BusinessObjects Business Intelligence platform Administrator Guide</i>.
User Name	The user name to access the data source if Authentication Mode is Use specified user name and password .
Password	The password to access the data source if Authentication Mode is Use specified user name and password .
Host Name	The host name of the server where the SAS/SHARE is running.
Port	The port to which to connect.
Catalog	<p>The name used to identify the connection to the query server.</p> <div>  Note A default catalog name is registered automatically with the query server the first time the connection is added to any multisource-enabled data foundation. </div>

9.3.3.3.1 Setting SAS data sets

Context

When you create an SAS connection, the [Set SAS Data Sets](#) dialog box lets you configure the data foundation to access multiple data sets that are not pre-defined to the SAS/SHARE server. These are data sets that are not included in the current SAS configuration.

Procedure

1. Select the **Use data sets that are non pre-defined to the SAS/SHARE server** option.
2. Click **Add** and, in the **Location** field, enter the path to the data set in the format required for the operating system that you are using.
3. In the **Library Name** field, enter a name to use to refer to the data set.
4. Click **Add** to add other data sets if needed.
5. Click **Finish**.

9.3.3.4 Login parameters for Oracle EBS connections

The following login parameters apply to Oracle EBS connections.

Parameter	Description
User Name	The user name to access the Oracle database server.
Password	The password to access the Oracle database server.
Service	The Oracle service name.
Authentication Mode	<p>The method used to authenticate the user's login credentials when accessing the EBS application:</p> <ul style="list-style-type: none">• Use specified user name and password: Uses the Oracle EBS User and Oracle EBS Password parameters defined for the connection.• Use Single Sign On: Uses the credentials associated with the user's account defined on the Central Management Server (CMS) when users log into the SAP BusinessObjects BI platform using Oracle EBS user name and password. For more information, see the <i>SAP BusinessObjects Business Intelligence platform Administrator's Guide</i>.
Oracle EBS User	The user name to access the application if Authentication Mode is Use specified user name and password .
Oracle EBS Password	The password to access the application if Authentication Mode is Use specified user name and password .
Language	The application language.
Application	The application name.
Security Group	The Oracle security group.

9.3.3.5 Parameters for OData connections


The following parameters apply to connections to OData data sources.

Login parameters for OData connections



Parameter	Description
Service Root URI	The URI string of the OData service. For example: <code>http://services.odata.org/OData/OData.svc</code>
Authentication Mode	The method used to authenticate the user's login credentials when accessing the data source: <ul style="list-style-type: none">• Use specified user name and password: Uses the User Name and Password parameters defined for the connection.
User Name	The optional user name for HTTP authentication.
Password	The optional password for HTTP authentication.
Proxy Address	The path to the HTTP proxy server (<host:port>). For example: <code>myproxy.com:8080</code>
Proxy User Name	The user name used to access the proxy server.
Proxy Password	The password used to access the proxy server.

Extended parameters for OData connections

Parameter	Description
Custom Authentication Parameters	Custom parameters used for authentication. They are attached to the URI but not traced to avoid exposing secured information. For example: <code>apikey=1234&authinfo=1234</code>
Column Selection	If selected, the Odata service provider executes the corresponding operation of the SQL query. If unselected, the OData driver executes the operation.
Supported Filter Conditions	
Sorting	

Parameter	Description
	<p> Note</p> <p>SAP recommends you do not use the data access driver to execute these operations, because it may affect the connection performance. Use it only if the service provider does not support or partially supports the operation.</p>

Configuration parameters for Odata connections


Parameter	Description
Connection Pool Mode	If using a connection pool, the method to use to keep the connection active.
Pool Timeout	If the Connection Pool Mode is set to Keep the connection active for , the length of time in minutes to keep the connection open.
Connection Timeout	<p> Restriction</p> <p>Specific to HTTP connections to OData and Web Service data sources.</p> <p>The time in seconds a connection remains active in case of no response from the data source. Default value is 10.</p> <p>The connection remains active indefinitely if Connection Timeout is set to 0.</p>
Cache Metamodel	<p>If selected, caches the metamodel in the connection so that the model does not need to be parsed and recreated for each connection call.</p> <p> Note</p> <p>SAP recommends you do not use the data access driver to execute this operation, because it may affect the connection performance.</p>

9.3.3.6 Parameters for XML and Web Services connections

The following parameters apply to connections to XML data sources and web services.

Parameter	Description
Location Type Protocol	If Location Type is set to Local , the protocol and login credential parameters are grayed.



Parameter	Description
	<p>If Location Type is set to Remote, select a Protocol to make the appropriate login parameters available.</p> <div> <p>i Note</p> <p>Schema files can be local even if Location Type is set to Remote. If the schema file is remote, the Protocol applies to both the data source and schema files.</p> </div>
Filepath or Pattern	<p>The path to a single XML file or the path to a folder containing multiple XML files. Files can be local or remote (HTTP, FTP, and SMB). If remote, the data-source is the location URL. Paths with MS Windows or UNIX styles are valid. Wildcards can be used. Blank characters must be replaced with %20.</p> <p>For example:</p> <ul style="list-style-type: none"> • C:\report.xml for a single file • C:\XMLFiles\ or C:\XMLFiles*.xml for multiple files • /home/user/xmlfiles/report.xml for a single file located on a UNIX machine • Remote locations: <ul style="list-style-type: none"> ◦ http://host:port/path/file ◦ ftp://host:port/path/file ◦ smb://server:port/path/file
Choose the XML schema	<p>If set to Explicitly indicate the XML Schema (XSD), the data access driver uses the XML schema that you enter in Schema File.</p> <p>If set to The XML schema is included in the XML files, the data access driver uses the XML schema included in the XML files.</p>
Schema File	<p>The path to the XML schema.</p> <p>Required if Choose the XML schema is set to Explicitly indicate the XML schema (XSD).</p>
User Name	The user name to access the XML files in remote connection.
Password	The password to access the XML files in remote connection.
SMB Domain	The domain used for SMB connections.
Proxy Address	<p>The path to the HTTP proxy server (<host:port>).</p> <p>For example: myproxy.com:8080</p>
Proxy User Name	The user name used to access the proxy server.
Proxy Password	The password used to access the proxy server.




Parameter	Description
Merge Files into One	<p>This Boolean value indicates if tables are automatically concatenated when a pattern is given as data source.</p> <p>For example, if the data source pattern is <code>report_*.xml</code>, then the driver concatenates all the tables from XML files that match the pattern.</p> <div>  Caution XML files must have the same structure. </div>
Web Service URL	The path to the web service over HTTP or HTTPS.

9.3.3.7 Configuration parameters for relational connections

The **Configuration Parameters** dialog box contains parameters that you can set to override default configuration options.

The following configuration parameters apply to most relational connections.

Parameter	Description
Connection Pool Mode	If using a connection pool, the method to use to keep the connection active.
Pool Timeout	If the Connection Pool Mode is set to Keep the connection active for , the length of time in minutes to keep the connection open.
Connection Timeout	<div>  Restriction Specific to HTTP connections to OData and Web Service data sources. </div> <p>The time in seconds a connection remains active in case of no response from the data source. Default value is 10.</p> <p>The connection remains active indefinitely if Connection Timeout is set to 0.</p>
Array Fetch Size	<p>The maximum number of rows authorized with each fetch from the database.</p> <p>For example, if you enter 20, and your query returns 100 rows, the connection retrieves the data in 5 fetches of 20 rows each.</p> <p>To deactivate array fetch, enter an Array Fetch Size of 1. The Data is retrieved row by row.</p> <div>  Note Deactivating array fetch size can increase the efficiency of retrieving your data, but slows server performance. The greater the value in the Array </div>

Parameter	Description
	Fetch Size , the faster your rows are retrieved. You must, however, ensure you have adequate client system memory.
Array Bind Size	This parameter is not used for universes created using the information design tool.
Login Timeout	The number of seconds before a connection attempt times out and an error message is displayed.
Query Timeout	<p> Restriction</p> <p>Specific to connections to Oracle data sources using JDBC middleware.</p> <p>The number of seconds before a query running on the database times out and is forced to end.</p>
Add File(s)	<p> Restriction</p> <p>Specific to connections to Apache Hadoop HIVE data sources.</p> <p>The paths to external resources to be added to the Hadoop Distributed Cache of the cluster. Typically, resources can be Python transform script files that you make available at query execution time. This parameter corresponds to the <code>add FILE HIVE</code> command line.</p> <p>You can define the paths to more than one file, separated by semi-colons. For example:</p> <pre>/tmp/foo.py;/tmp/bar.py</pre> <p>Only UNIX style paths are valid.</p>
Add Jar(s)	<p> Restriction</p> <p>Specific to connections to Apache Hadoop HIVE data sources.</p> <p>The paths to external JAR files to be added to the Java classpath. This parameter corresponds to the <code>add JAR HIVE</code> command line.</p> <p>You can define the paths to more than one JAR file, separated by semi-colons. For example:</p> <pre>/usr/lib/hive/hive-contrib-1.jar;/usr/lib/hive/hive-contrib-2.jar</pre> <p>Only UNIX style paths are valid.</p>
JDBC Driver Properties (key=value,key=value)	Values for JDBC driver properties. You can define the value of more than one property, separated by commas. For example, the following value for JDBC

Parameter	Description
	<p>Driver Properties sets the <code>oracle.jdbc.defaultNChar</code> and <code>defaultNChar</code> driver properties:</p> <pre>oracle.jdbc.defaultNChar=true,defaultNChar=true</pre> <p>Note</p> <p>If a property is defined in the <driver>.sbo file, the value defined in this parameter is used. For more information about SBO files, see the <i>Data Access Guide</i>.</p>
Owner Name	For DB2 connections, this parameter adds the name of the owner of the table as a prefix on the table name, to match the DB2 convention for naming tables.
Table Suffix	For DB2 connections, this parameter adds a suffix on the table name, to match the DB2 convention for naming tables.


9.3.3.8 Custom parameters for relational connections




The Custom Parameters dialog box lets you override the value of certain parameters. You can also add parameters and their values.

Parameter	Description
ConnectInit	The value is added to the SQL and is run once when a user connects to the database.
Hint	For Oracle connections, the value is used by the Oracle query optimizer to choose an execution plan. Consult your Oracle documentation for full information on the Hints that can be used, and how they can be used to optimize queries.

9.3.3.9 Login and schema parameters for CSV file connections

The following parameters apply to connections to Comma-Separated Value (CSV) files. See the related topic for more information about file format and regional settings for CSV file connections.

Parameter	Description
Location Type Protocol	<p>The Location Type and Protocol parameters guide you through setting the necessary parameters for your data source.</p> <p>If Location Type is set to Local, the protocol and login credential parameters are grayed.</p> <p>If Location Type is set to Remote, select a Protocol to make the appropriate login parameters available.</p> <div>  Note </div> <p>Schema files can be local even if Location Type is set to Remote. If the schema file is remote, the Protocol applies to both the data source and schema files.</p>
Datasource	<p>Filepath or Pattern</p> <p>The path to a single CSV file or the path to a folder containing multiple CSV files. Files can be local or remote (HTTP, FTP, and SMB). If remote, the data-source is the location URL. Paths with MS Windows or UNIX styles are valid. Wildcards can be used. Blank characters must be replaced with %20.</p> <p>For example:</p> <ul style="list-style-type: none"> • C:\report.csv for a single file • C:\CSVFiles\ or C:\CSVFiles*.csv for multiple files • /home/user/csvfiles/report.csv for a single file located on a UNIX machine • Remote locations: <ul style="list-style-type: none"> ◦ http://host:port/path/file ◦ ftp://host:port/path/file ◦ smb://server:port/path/file
Schema Detection	<p>The method used to detect the schema of the CSV file. Possible values are:</p> <ul style="list-style-type: none"> • auto The data access driver finds the schema automatically. File Type must be set to delimited. • no detection The data access driver skips comment lines, analyzes the first line, and determines the number of columns, but not the column types. File Type must be set to delimited. • ddl The data access driver uses a Data Definition Language (DDL) file to detect the schema. • sqlddl The data access driver uses a DDL file corresponding to the Standard SQL to detect the schema.

Parameter	Description
Schema File	<p>The path to one DDL or SQLDDL schema file. If you want to define schemas for several tables, use an SQLDDL file.</p> <p>Required if Schema detection is set to ddl or sqlddl.</p> <p>If the schema file is remote, the protocol of the file must match the Protocol parameter setting.</p>
Probe Rows	<p>The method to parse lines for checking column information (name, type, size, and nullability). Required if Schema Detection is set to auto. Possible values are:</p> <ul style="list-style-type: none"> auto The driver parses the file until a type is detected for each column. If there is no NULL value in the first row, the parsing terminates after the first row. <div>  Caution This method may lead to type conversion conflicts in the case where only the first rows are parsed and other types are used in the next rows. </div> <ul style="list-style-type: none"> int Parses a specific number of rows. This setting can be used as a trade-off between scalability for big CSV files and a low data quality of CSV files. <div>  Caution This method may lead to conflicts if other types are used in the rows that are not parsed. </div> <ul style="list-style-type: none"> all Parses the whole file. This method allows you to find the longest string value, which corresponds to the column size for non-numeric values. <div>  Note The recommended method is all to allow for correct detection of column information. Since this is the slowest detection method, use DDL files if CSV files are large. </div>
Number of Probe Rows	<p>The number of CSV file lines that are parsed to check the column types.</p> <p>Required if Probe Rows is set to int.</p>
Credentials	<p>The User Name and Password to access the CSV files in a remote connection.</p> <p>If Protocol is set to SMB (Windows Share), enter the SMB Domain for the connection.</p>

Parameter	Description
HTTP Proxy	<p>If Protocol is set to HTTP, enter the proxy parameters for the connection.</p> <p>Proxy Address: The path to the HTTP or FTP proxy server (<host:port>). For example: <code>myproxy.com:8080</code></p> <p>Proxy User Name: The user name used to access the proxy server.</p> <p>Proxy Password: The password used to access the proxy server.</p>


Related Information



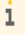

[File format and regional settings for CSV file connections](#) [page 121]

9.3.3.10 File format and regional settings for CSV file connections

The following parameters apply to connections to Comma-Separated Value (CSV) files.

File Settings

Parameter	Description
File Charset	<p>The character set used in the CSV files.</p> <div>  Note All files must have the same character set. </div>
File Type	<p>The file type can be one of the following:</p> <ul style="list-style-type: none"> • delimited The CSV file entries are delimited by a separator. • fixed The CSV file entries have fixed width.
Lenient Mode	<p>If selected, invalid rows are skipped automatically (not enough or too many columns).</p>
Column Names in First Line	<p>This Boolean value tells if the first row of the CSV file contains column names.</p>

Parameter	Description
	<p> Note</p> <p>If the file does not contain any column names and if schema detection is enabled, then the data access driver names columns as <code>col1</code>, <code>col2</code>, ..., <code>col<n></code>.</p>
Merge Files into One	<p>This Boolean value indicates if tables are automatically concatenated when a pattern is given as data source.</p> <p>For example, if the data source pattern is <code>report_*.csv</code>, then the driver concatenates all the tables from CSV files that match the pattern.</p> <p> Caution</p> <p>CSV files must have the same structure.</p>
Number of Comment Lines in the Beginning	<p>The number of rows that contain comments at the beginning of the CSV file. Maximum is 1000.</p>
Separator	<p>A character that is used to separate CSV file entries. It must be different from text qualifier and escape character.</p> <p> Note</p> <p>If the tab key is used to separate entries, the word TAB can be set as separator.</p>
Text Qualifier	<p>A character that surrounds a file entry, for example quote (') or double quote (").</p> <p>If you do not want to use a text qualifier, use a character that is not used in the CSV file to prevent the data access driver from using the default value.</p>
Escape Character	<p>A character that allows the text qualifier to be treated as literal text.</p> <p> Note</p> <p>Text qualifier and escape character must be different.</p>

Regional Settings

Parameter	Description
Decimal Separator	Default value is period (.). For example: 100.20.
Thousands Grouping Character	Default value is comma (,). For example: 1,000.20.
Date Format	Date and time formats used in the CSV files. They must match the format in the CSV files to enable the driver to recognize the date and time formats and to parse them. Default values are: <ul style="list-style-type: none">• <code>yyyy-MM-dd</code> for dates• <code>yyyy-MM-dd HH:mm:ss</code> for timestamps• <code>HH:mm:ss</code> for times
Timestamp Format	
Time Format	

9.4 Creating an OLAP connection

Prerequisites

You use the New OLAP Connection wizard to create local and secured connections to an OLAP data source.

Before you create a local connection in the information design tool, you must have a project available in the Local Projects View. For more information on creating local projects, see the related topic.

Note

The OLAP connections you create in the information design tool are not supported by the universe design tool. Also, OLAP connections created in the universe design tool are not available for building universes in the information design tool.

Procedure

1. Do one of the following:
 - To create a local connection, select the project folder in the Local Projects View. Select **File > New > OLAP Connection**.
 - To create a secured connection, in the Repository Resources View, open a session on the repository where you want to create the secured connection. Right-click the Connections folder or subfolder in the repository, and select **Insert OLAP Connection**.

2. Follow the steps in the New OLAP Connection wizard to enter the following information:

- Name for the connection
- Middleware driver for the target database
- Authentication parameters to connect to the OLAP data source
- OLAP cube to connect to

If you need help on a particular step, click the help icon in the wizard dialog box.

Related Information

[Creating a local project](#) [page 79]

[Select an OLAP middleware driver](#) [page 124]

[About connections](#) [page 99]

9.4.1 Select an OLAP middleware driver

This section describes the OLAP driver selection page of the New OLAP Connection wizard.

You select an OLAP driver to connect to the OLAP server. The OLAP driver maps information from the OLAP server middleware to the user interface of the SAP BusinessObjects application.

Depending on your target OLAP server, expand the middleware node, and select the target driver.

i Note

If you are using the information design tool from a Crystal Server 2011 installation, SAP middleware drivers are not available.

9.4.2 Set login parameters for OLAP data sources

The connection parameters vary depending on the type of data source for which you are defining the connection. Select from the related topics the link to more information about the connection parameters.

Related Information

[Login parameters for OLAP connections](#) [page 125]

[Login parameters for SAP NetWeaver BW and ERP connections](#) [page 107]

9.4.2.1 Login parameters for OLAP connections

The following parameters apply to most OLAP connections.

For a description of login parameters for SAP NetWeaver BW (BICS Client), see the related topic.

Login parameter	Description
Authentication mode	<p>The method used to authenticate the user's login credentials when accessing the data source:</p> <ul style="list-style-type: none">• Use specified user name and password: Uses the User Name and Password parameters defined for the connection.• Use BusinessObjects credential mapping: Uses the database credentials associated with the user's account defined on the Central Management Server (CMS) to connect to the data source. The database credentials are set in the User's Properties in the Central Management Console. For more information, see the <i>SAP BusinessObjects Business Intelligence platform Administrator's Guide</i>.• Use Single Sign On: This authentication mode is used to support end-to-end single sign on defined in the Central Management Server (CMS). If you use an external authentication source (for example LDAP), the CMS and the data source must be configured to use this external authentication source. For more information on single sign on, see the <i>SAP BusinessObjects Business Intelligence platform Administrator Guide</i>.
Host Name	<p>This parameter applies only to SAP HANA connections.</p> <p>The name of the server hosting the data source. Do not include the port number.</p>
Instance Number	<p>This parameter applies only to SAP HANA connections.</p> <p>The SAP HANA instance number, which represents the second and third digits of the port number. Select a number between 0 and 99. For example, if the port number is 30215, the instance number is 2.</p>
Server	<p>For MSAS connections, the URL path, for example:</p> <p><a href="http://<server_name>/olap_2005/msmdpump.dll">http://<server_name>/olap_2005/msmdpump.dll</p> <p>For Essbase connections, the server name for the data source.</p>
User Name	<p>The user name to use to access the OLAP server if the Authentication mode is Use specified user name and password.</p>
Password	<p>The password to use to access the OLAP server if the Authentication mode is Use specified user name and password.</p>
Language	<p>The language that will be used for the connection.</p>
Auto Reconnect	<p>This parameter applies only to SAP HANA connections.</p>

Login parameter	Description
	If selected, the application automatically reconnects to the host server if the connection fails.
Use SSL	This parameter applies only to SAP HANA connections. If selected, uses SSL protocol to connect to the host server.
Fetch Size	This parameter applies only to SAP HANA connections. The maximum number of rows authorized with each fetch from the database. The recommended Fetch Size for OLAP connections to SAP HANA is 7000.

Related Information

[Login parameters for SAP NetWeaver BW and ERP connections](#) [page 107]

9.4.2.2 Login parameters for SAP NetWeaver BW and ERP connections

The following parameters apply to connections to SAP NetWeaver BW (relational and BICS Client connections), and SAP ERP.

To set ABAP function and InfoSet parameters for SAP ERP connections, after entering the login parameters, click **Next**.

Parameter	Description
Authentication Mode	<p>The method used to authenticate the user's login credentials when accessing the data source:</p> <ul style="list-style-type: none"> • Use specified user name and password: Uses the User Name and Password parameters defined for the connection. • Use BusinessObjects credential mapping: Uses the database credentials associated with the user's account defined on the Central Management Server (CMS) to connect to the data source. The database credentials are set in the User's Properties in the Central Management Console. For more information, see the <i>SAP BusinessObjects Business Intelligence platform Administrator's Guide</i>. • Use Single Sign On: This authentication mode is used to support end-to-end single sign on defined in the Central Management Server (CMS). If you use an external authentication source (for example LDAP), the CMS and the data source must be configured to use this external authentication source.

Parameter	Description
	For more information on single sign on, see the <i>SAP BusinessObjects Business Intelligence platform Administrator Guide</i>
Client Number	The number used to identify the client on the SAP system.
User Name	The user name to access the data source if Authentication Mode is Use specified user name and password .
Password	The password to access the data source if Authentication Mode is Use specified user name and password .
Language	<p>The 2-character ISO language code for the language to be used for the connection to the data source. For example, EN for English.</p> <div> <p>i Note</p> <p>In some cases, select the language from the list.</p> </div>
Save Language	<p>Specifies which language to use for the connection:</p> <ul style="list-style-type: none"> • If you select the Save Language option, the value from the Language parameter is used. • If you clear Save Language, the value from the user's session (Preferred Viewing Locale) is used.
System ID	<p>The 3-character SAP system ID.</p> <div> <p>i Note</p> <p>Required for both application and message server types.</p> </div> <div> <p>i Note</p> <p>For a successful connection to the message server, you need to add the message server system ID to the following file on the machine hosting the application:</p> <p>C:\WINDOWS\system32\drivers\etc\services</p> <p>At the end of the existing file, add the line:</p> <p>sapmsXXX <tab> 3601/tcp</p> <p>where sapms means SAP message server, xxx is the is the System ID of the server that is used, and 3601/tcp is the default TCP port used for communication.</p> </div>
Server type	<ul style="list-style-type: none"> • Select Application Server to connect directly to the SAP server without using load balancing. • Select Message Server to benefit from SAP load balancing capabilities.

Parameter	Description
Server Name for Application Server	The name of the SAP application server.
System Number for Application Server	The system number of the SAP application server. This is a two-digit integer between 00 and 99.
Server Name for Message Server	The name or IP address of the SAP message server used for load balancing.
Group Name for Message Server	Name of the Logon group; a set of dedicated application servers used to log on.

The following parameters apply only to SAP NetWeaver BW connections:

Parameter	Description
Use Custom ID Program Mapping	<p>An optional parameter for SAP NetWeaver BW relational connections only.</p> <p>The Program ID Mapping defines the program IDs for the callback that SAP NetWeaver BW uses to contact the data federation server. Enter Program ID Mapping as a list of one or more server name=program ID pairs separated by the semi-colon character (;). For example:</p> <p><code><MySIA.DF_Server1>=RFC1;<MySIA.DF_Server2>=RFC2</code></p> <p>Each program ID must match the name of an RFC destination created on SAP NetWeaver BW.</p> <p>If this parameter is not defined, the data federation server automatically creates an RFC destination.</p> <p>For more detailed information, see the description of the programIDMapping connector property in the Data Federation Administration Tool Guide.</p>
Use Custom Gateway	<p>An optional parameter for SAP NetWeaver BW relational connections only.</p> <p>In Gateway Host Name, enter the name of the server hosting the SAP NetWeaver BW gateway.</p> <p>In Gateway Service Name, enter the name or port number of the SAP NetWeaver BW gateway service.</p> <p>If this option is not selected, SAP NetWeaver BW provides the gateway host name and service name via an RCF.</p>
InfoProvider	For SAP NetWeaver BW relational connections, the name of the InfoCube or MultiProvider to be used as the fact table in the center of the snowflake schema in the data foundation.
Catalog	For SAP NetWeaver BW relational connections, the name used to identify the connection to the query server.

Parameter	Description
	<p>i Note</p> <p>A default catalog name is registered automatically with the query server the first time the connection is added to any multisource-enabled data foundation.</p>

9.4.3 Select an OLAP cube

The following options apply to associating a cube with the OLAP connection.

i Note

For connections using the **SAP BICS Client**, select the **Specify a cube in the connection** option. Open a catalog, then select the BEx Query for the connection.

Option	Description
Do not specify a cube in the connection	Select this option to create the connection without specifying a cube. In this case, each time you access the connection, either when building a business layer or in a query and reporting tool, you will be prompted to select a cube.
Specify a cube in the connection	<p>Select this option to always associate a cube with the connection.</p> <p>The cube selection page lists the cubes available for the target database. You can enter a search string in the search text box. Select the cube in the list.</p>

9.5 Creating a connection shortcut

Context

When you publish a connection, you have the option of creating a connection shortcut in the Local Projects View. Use the following procedure to create a connection shortcut for an existing secured connection.

You must have a local project in the Local Projects View.

Procedure

1. In the Repository Resources View, open a session on the repository where the secured connection is stored.

2. In the Connections folder or subfolder, right-click the connection name.
 - For OLAP connections, select **Create OLAP Connection Shortcut**.
 - For relational connections, select **Create Relational Connection Shortcut**
3. In the [Select a Local Project](#) dialog box, select the project in which you want to create the shortcut.

Related Information

[Creating a local project](#) [page 79]

[About connection shortcuts](#) [page 101]

9.6 Editing local and secured connections

Procedure

1. To open the connection in the editor, do one of the following:

Option	Description
To open a local connection	Double-click the connection name in the Local Projects View.
To open a secured connection	In the Repository Resources View, open a session on the repository where the connection is published. In the Connections folder or subfolder, double-click the connection name.

2. To edit the connection name or description, click the **General Information** tab.
3. To edit the connection parameters, click **Edit**.

For local connections, you can also right-click the connection name in the Local Projects View and select **Edit Connection**.

4. To change the middleware driver for relational connections, select **Change Driver**. Select a new driver and enter the new connection parameters.
5. To test the availability of the database server, click **Test Connection**.

You can also right-click the connection or shortcut name in the Local Projects View and select **Test Connection**.

6. Save the connection information by clicking the Save icon in the main tool bar.

Related Information

[Showing values in a relational connection](#) [page 131]

[Showing values in an OLAP connection](#) [page 132]

9.7 Editing connection shortcuts

Context

You can edit the name and description of a connection shortcut. You can also change the shortcut to reference a different connection in the same repository where the existing connection is published.

Procedure

1. Open the connection shortcut in the editor by double-clicking the shortcut name in the Local Projects View.
2. You can enter or change text in **Shortcut Name** and **Description**.
3. To change the connection that the shortcut references, click **Change Connection**.

You can also right-click the shortcut name in the Local Projects View and select **Change Connection**.

4. To test the referenced connection, click **Test Connection**.

You can also right-click the shortcut name in the Local Projects View and select **Test Connection**.

5. Save the shortcut by clicking the Save icon in the main tool bar.

Related Information

[About connection shortcuts](#) [page 101]

9.8 Showing values in a relational connection

Procedure

1. Open the connection in the editor:

Option	Description
To open a local connection	Double-click the connection name in the Local Projects View.
To open a secured connection	In the Repository Resources View, open a session on the repository where the connection is published.

Option	Description
	In the Connections folder or subfolder, double-click the connection name.

- Click the **Show Values** tab.
- In the **Catalog** pane, double-click the table name (to show all columns), or double-click a column name.
To see what you can do in the pane where values are displayed, see the related topic.

Related Information

[Showing values in a data source](#) [page 185]

9.9 Showing values in an OLAP connection

Procedure

- Open the connection in the editor:

Option	Description
To open a local connection	Double-click the connection name in the Local Projects View.
To open a secured connection	In the Repository Resources View, open a session on the repository where the connection is published. In the Connections folder or subfolder, double-click the connection name.

- To browse the objects in the cube and their properties, click the **Browse Metadata** tab.
Select an object to display its properties in the properties pane.
- To run an MDX query on the cube, select the **Query** tab.

Note

MDX querying is not available for OLAP connections used for direct access (such as direct access to a BEx Query or SAP HANA information model).

- Build an MDX query in the **MDX Query** pane by dragging and dropping objects from the **OLAP Metadata** pane and entering MDX statements.
- To validate the MDX, click **Parse**.
- To run the query, click **Run**.

10 Working with data foundations

10.1 About data foundations

A data foundation contains a schema of relevant tables and joins from one or more relational databases that are used as a basis for one or more business layers.

You reference relational connections in the data foundation. You insert tables and joins from the databases referenced in the connections.

Using the Data Foundation Editor, you can enhance the data foundation by adding federated tables (designed in the federation layer), derived tables, alias tables, calculated columns, additional joins, contexts, prompts, and lists of values. The availability of some features depends on the data foundation type. See the related topic for more information about data foundation types.

You can build any number of business layers on the same data foundation. In this case the data foundation becomes the basis for multiple universes.

Related Information

[About data foundation types](#) [page 133]

[About the federation layer](#) [page 196]

[How to build a data foundation](#) [page 138]

10.1.1 About data foundation types

Single-source and multisource-enabled are two types of data foundations that allow you to take advantage of different data foundation features.

Related Information

[About single-source data foundations](#) [page 133]

[About multisource-enabled data foundations](#) [page 134]

10.1.2 About single-source data foundations

Single-source data foundations support a single connection. The connection can be local or secured, which means you can publish universes based on the data foundation either locally or to a repository.

Single-source data foundations support database-specific SQL syntax for derived tables, calculated columns, and join expressions. Database-specific SQL syntax allows functions or operators that are offered by a specific database and not by standard SQL-92 (for example, Oracle analytical functions). You must select single-source if you want to publish to a local folder the universes that are based on this data foundation.

Single-source data foundations are recommended for the following situations:

- You want to work exclusively with database-specific SQL syntax.
- You want to publish the universe locally and work outside of a repository.

Related Information

[About multisource-enabled data foundations](#) [page 134]

[How to build a data foundation](#) [page 138]

[Changing a connection in a data foundation](#) [page 143]

10.1.3 About multisource-enabled data foundations

Multisource-enabled data foundations support one or more connections. You can add connections when you create the data foundation and anytime later. Multisource-enabled data foundations only support secured connections, and universes based on this type of data foundation can only be published to a repository.

Multisource-enabled data foundations support most relational connections supported in single-source data foundations. In addition, multisource-enabled data foundations support the following relational connections that are not supported in single-source data foundations:

- SAP NetWeaver BW connections
- SAS connections

Connections for multisource-enabled data foundations are managed by the data federation service. For information about tuning the data federation service, see the *Data Federation Administration Tool Guide*.

The federation layer is available in multisource-enabled data foundations. It allows you to create federated tables that you can then include in the data foundation.

SQL-92 standard syntax is the default for calculated columns, derived tables, and join expressions. In addition, the SAP BusinessObjects SQL database functions are available. You can use database-specific SQL syntax in a multisource-enabled data foundation by defining a database-specific derived table or calculated column. Database-specific SQL syntax allows functions or operators that are offered by a specific database and not by standard SQL-92 (for example, Oracle analytical functions).

i Note

For database functions, the SAP BusinessObjects syntax can be different from the syntax of the same function provided by database-specific SQL.

Multisource-enabled data foundations are required in the following situations:

- You want to insert tables and joins from more than one relational data source, or create federated tables.

- You want to use SAP NetWeaver BW or SAS connections.
- You want to use SQL-92 standard syntax and SAP BusinessObjects SQL functions.

For more information on these situations, see the related topics.

Related Information

[Data foundations with multiple connections](#) [page 135]

[About the federation layer](#) [page 196]

[SQL expressions in multisource-enabled data foundations](#) [page 136]

[About single-source data foundations](#) [page 133]

[How to build a data foundation](#) [page 138]

[Using SAP NetWeaver BW data sources](#) [page 40]

10.1.3.1 Data foundations with multiple connections

To be able to add multiple connections to a data foundation, you must select the multisource-enabled type when you create the data foundation.

You can select multiple connections when you create the data foundation. You can also add connections to an existing multisource-enabled data foundation. Connections must be secured, and therefore available in a repository. The connections are represented by a connection shortcut in the local project.

The connections in a multisource-enabled data foundation have the following additional properties:

- A short name used to identify the connection in the data foundation and to modify the table name in SQL expressions. You specify the short name when adding the connection. This name must be unique within the data foundation and is limited to forty characters. If you change the short name for the connection, the SQL expressions are automatically updated with the new name.
- A color for the connection. This color is used in the table header in data foundation views. You select the color when adding the connection. You can change the color for a connection at any time.
- A catalog used to identify the connection to the query server. A default catalog name is registered automatically with the query server the first time the connection is added to any multisource-enabled data foundation.
- For SAP NetWeaver BW connections, properties related to the automatic insertion of tables and joins. For more information on these properties, see the related topic.

In a multisource-enabled data foundation, the table name as it appears in SQL expressions has the format:

`@catalog(short name)."database_qualifier.database_owner"."table_name"`

A multi-source join can be created between tables from different connections. You can use the **Detect Joins** command to detect joins between tables referenced in different connections, or explicitly define them with the **Insert Join** command.

Related Information

[Using SAP NetWeaver BW data sources](#) [page 40]

[About connections in the data foundation](#) [page 141]

[Changing a connection in a data foundation](#) [page 143]

10.1.3.2 SQL expressions in multisource-enabled data foundations

SQL expressions that define joins, calculated columns, and derived tables in a multisource-enabled data foundation use SQL-92 ANSI standard syntax.

In SQL-92 expressions, you can include SAP BusinessObjects database functions. The SQL syntax can be different from the syntax of the same function provided by database-specific SQL. For more information, see the related topic.

In SQL-92 expressions, you can include @functions. Which @functions you can include depends on the type of expression. For more information, see the related topic.

In order to use functions or operators that are offered by the database and not by SQL-92 (for example, Oracle analytical functions), you define database-specific calculated columns and derived tables. An option in the SQL Expression Editor allows you to use database-specific SQL.

Database-specific calculated columns and derived tables support the SQL syntax of the associated connection. The following rules apply to database-specific SQL expressions:

- You can reference only standard tables and database-specific derived tables in a single connection.
- You cannot reference tables in SAS or SAP NetWeaver BW connections.
- You can include @functions with certain restrictions. For more information, see the related topic.

Related Information

[SAP BusinessObjects SQL function reference for multisource-enabled universes](#) [page 360]

[About @Functions](#) [page 423]



10.2 About the Data Foundation Editor

This topic describes how to navigate the Data Foundation Editor. For steps to help you build the structure of your data foundation, see [How to build a data foundation](#) [page 138].

The Data Foundation Editor is divided into a data foundation view pane, a properties pane, and browsing panes.

The data foundation view is a graphical representation of the tables and joins. The **Master** view contains all tables and joins and cannot be deleted. You can define custom views that contain subsets of the tables. Access the views by the tabs at the bottom of the view pane. For more information about custom views, see the related topic.

The properties pane displays the properties of the data foundation object that is currently selected (the entire data foundation, a table, a column, or a join). To edit properties that apply to the entire data foundation, see the related topic.

In the data foundation view, you can work on tables and joins using commands in the **Insert**  and **Detect**  menus, or by clicking objects directly in the view.

The browsing panes allow you to work with different elements of the data foundation. Access the panes by clicking the corresponding tab:

- **Connections**
- **Data Foundation** (displays a tree view of the tables and joins)
- **Aliases and Contexts**
- **Parameters and Lists of Values**
- **Federation Layer**

For more information about what you can do in each of the browsing panes, see the related topic.


Navigating the data foundation view

To access a menu of commands that are available on tables, right-click the table header in the data foundation view. To select multiple tables, click the table headers while holding down the **CTRL** key.

To access commands that are available on columns, right-click the column name in the table in the data foundation view.

Several commands in the table right-click menu are available to help you locate related tables in the data foundation:

- **Select Related Tables** automatically selects all tables linked by joins to the selected table.
- **Highlight Related Tables** grays any table that is not linked to the selected table by a join.
- **Highlight Aliases** grays all tables except the selected original table and its alias tables.
- **Highlight Original Table** grays all tables except the selected alias table and the original table it is based on.
- **Center on Selection** lets you temporarily change the zoom on the data foundation display so that all the tables in a selection are visible in the display window.

You can use the search panel to perform advanced searches on the data foundation. To open the search panel, click .

For information on commands you can use to change the display of objects in the data foundation view, see the related topic.

Related Information

[Inserting a custom data foundation view](#) [page 190]

[About connections in the data foundation](#) [page 141]

[About contexts](#) [page 175]

[About parameters and lists of values in the data foundation](#) [page 180]

[About the federation layer](#) [page 196]

[About data foundation properties](#) [page 180]

[Searching for tables and columns in the data foundation](#) [page 191]

[Centering the view on a selection](#) [page 192]

[Changing the display of objects in the data foundation](#) [page 193]

10.3 How to build a data foundation

Prerequisites

Before you begin:

- You need a local project in which to create the data foundation.
- In the local project, you need the relational connection or connection shortcuts to secured relational connections. Multisource-enabled data foundations require connection shortcuts.

For links to more detailed information on each step, see the Related Topics.

Procedure

1. To start the [New Data Foundation](#) wizard, do one of the following:
 - Right-click a relational connection or connection shortcut in the Local Projects View and select ► **New** ► **Data Foundation** .
 - Right-click the project folder in the Local Projects View and select ► **New** ► **Data Foundation** .

The data foundation is created in a .dfx file in the local project. It opens automatically in the Data Foundation Editor.

2. If you want use federated tables in your data foundation (multisource-enabled data foundations only), create the federated tables in the **Federation Layer**.
3. In the Data Foundation Editor, insert tables into the data foundation:

Option	Command
To insert tables from the connection	In the Connection pane, open and browse the tables in the connection. You have tools to search for and filter the tables in the Connection pane. Drag and drop the tables you want into the data foundation Master view.
To insert tables using a wizard	<p>In the data foundation view tool bar, select ► Insert ► Insert Tables . You have the option to detect and insert keys, joins, cardinalities, and row counts automatically.</p> <div> <p>i Note</p> <p>To detect joins between tables referenced by different connections, you need to use the Detect Joins command.</p> </div>
To insert federated tables (multisource-enabled data foundations only)	In the data foundation view tool bar, select ► Insert ► Federated Table .

4. Insert joins:

Option	Command
To insert joins manually	<p>In the data foundation view, click the column name in the first table and drag it to the column in the second table. A join path appears between the two tables.</p> <p>You can also insert joins by opening the Edit Join dialog box. In the data foundation view tool bar, select ► Insert ► Insert Join .</p>
To detect joins	In the data foundation view tool bar, select ► Detect ► Detect Joins .

To edit a join, double-click the join path. For more information about editing and detecting joins, see the related topics.

5. Check the cardinality of the joins in the data foundation. In the data foundation view tool bar, select ► **Detect** ► **Detect Cardinalities** .

In the [Detect Cardinalities](#) dialog box, you can set or detect cardinality for any or all joins.

6. You can enhance the function of the data foundation in several ways, for example:

- Insert calculated columns
- Insert derived tables
- Insert alias tables
- Insert parameters with optional prompts
- Insert lists of values to be associated with a prompt
- Set SQL options and SQL generation parameters in the data foundation properties

7. Verify the join paths and resolve any loops. Use the commands in the **Aliases and Contexts** pane to detect aliases and contexts automatically.
8. Run an integrity check to validate the tables, columns, and joins in the data foundation. Right-click the data foundation name in the **Data Foundation** pane, and select **Check Integrity**.
9. Save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

Some commands to help you maintain the data foundation are listed below.

- If you change tables or columns, use **Show Local Dependencies** to find the business layers and objects that might be impacted by the changes.
- Refresh the structure to update the data foundation with changes that have occurred to the databases referenced in the connections.
- You can change a connection, or for multisource-enabled data foundations, add or remove a connection.
- You can enhance the ease of maintenance by creating custom views, grouping tables using families, arranging the table display, and inserting comments.

Related Information

[About data foundation types](#) [page 133]

[Creating a local project](#) [page 79]

[Creating a relational connection](#) [page 102]

[Creating a connection shortcut](#) [page 129]

[About the Data Foundation Editor](#) [page 136]

[Inserting tables into the data foundation](#) [page 151]

[Searching for tables in the Connection pane](#) [page 145]

[Filtering the tables in the connection by table type](#) [page 147]

[About the federation layer](#) [page 196]

[Inserting and editing a join](#) [page 161]

[Detecting and setting cardinalities](#) [page 165]

[Inserting a calculated column](#) [page 166]

[About derived tables](#) [page 168]

[About alias tables](#) [page 172]

[About parameters and lists of values in the data foundation](#) [page 180]

[About data foundation properties](#) [page 180]

[Resolving loops](#) [page 177]

[Running check integrity](#) [page 316]

[Showing local dependencies in the data foundation](#) [page 187]

[About refreshing a data foundation](#) [page 188]

[Changing a connection in a data foundation](#) [page 143]

[Adding connections to a data foundation](#) [page 143]




[Inserting a custom data foundation view](#) [page 190]
[Grouping tables using families](#) [page 194]
[Auto-arranging tables in the data foundation view](#) [page 193]
[Inserting a comment into the data foundation view](#) [page 192]

10.4 About connections in the data foundation



Connections in the data foundation are listed in the [Connection](#) pane of the Data Foundation Editor. Some connections allow for multiple databases (called qualifiers), with different owners:

- Some data sources provide both qualifiers and owners (for example, MS SQL Server)
- Some data sources only provide qualifiers (for example, MySQL and text files)
- Some data sources only provide owners (for example, Oracle, SAP HANA, DB2, and Teradata)

For single-source data foundations, the [Connection](#) pane display is as follows:

-  Qualifiers, if available, are listed under the connection
 -  Owners, if available, are listed under each qualifier
 -  Tables are listed under each owner
 - Columns are listed under each table

For multisource-enabled data foundations, the [Connections](#) pane display is as follows:

-  Qualifier.Owners, known as the schema, are listed under the connection (or owners only, if the data source does not provide qualifiers).
 -  Tables are listed under each schema
 - Columns are listed under each table

The list of tables is sorted alphabetically.


By default, tables are listed for all qualifiers and owners. To list only tables for the currently used qualifier/owner,

click the [Show Qualifiers and Owners](#) icon  to deselect it.


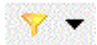
When the connection contains more than 1000 qualifier/owners or tables, the qualifier/owners or tables are grouped in packages of 1000. The packages are listed with the first several letters of the first and last qualifier/owner or table names in the package surrounded by brackets, for example:


 [AAAA....] - [MMMM]

 [NNNN...] - [ZZZZZ]

For tables that are already inserted into the data foundation, the table icon displays a green check mark: . You can insert a table into the data foundation by double-clicking the table name in the [Connections](#) pane.

Navigating the list of tables in the *Connection* pane

The **Show/Hide Table Search** icon,  lets you search for tables in a connection to obtain a filtered list of tables. Some connections have different table types (for example, in an SAP HANA connection, you can have several table types, including **Analytic View** and **Calculation View**). The **Filter by Table Type** icon  lets you select table types to filter the list of tables displayed in the connection.

For SAP HANA connections, the **Filter Information Models** icon  by default filters the list to show only tables representing information models.

The table search, filter by table type, and filter on information models can be used in combination. For more information on searching and filtering, see the related topics.

Operations on connections

You can do the following tasks on connections from the *Connection* pane:

- **Change** lets you change the connection and its associated properties. For more information about this task, see the related topic.
- **Open** opens the connection or connection shortcut properties in the Connection Editor.
- **Test** lets you test if the database referenced by the connection is available.

In addition, for multisource-enabled data foundations, you can do the following tasks from the *Connections* pane:

- **Add Connections** lets you add connections to the data foundation. For more information about this task, see the related topic.
- **Remove** lets you remove a connection from the data foundation. The connection itself remains in the repository and registered with the catalog name.

Related Information

[Searching for tables in the Connection pane](#) [page 145]

[About tables in the data foundation](#) [page 148]

[Filtering the tables in the connection by table type](#) [page 147]

[Filtering tables in an SAP HANA connection by information model](#) [page 148]

[Changing a connection in a data foundation](#) [page 143]

[Adding connections to a data foundation](#) [page 143]

[About the Connection Editor](#) [page 102]

[Setting connection display preferences for the Data Foundation Editor](#) [page 27]


10.4.1 Adding connections to a data foundation

Prerequisites

To add connections to the data foundation, the following conditions are required:

- The data foundation type must be multisource-enabled.
- The connections to add must be relational, secured connections.
- For each connection to add, you must create a connection shortcut in the local project where the data foundation is stored.

Procedure

1. Double-click the data foundation name in the local project to open the Data Foundation Editor.
2. In the Data Foundation Editor, click the **Connections** tab.
3. In the *Connections* pane, click the **Add Connections** icon .
The *Add Connections* dialog box lists the available connections, including the connections currently defined in the data foundation.
4. Select the connection shortcut name of each connection you want to add, and click **Next**.
5. In the *Connections Properties* dialog box, you can define additional connection properties. A dialog box opens for every connection added.
For more information about multisource-enabled connection properties, see the related topic.
6. When you have finished defining the properties for additional connections, click **Finish**.
7. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[Creating a connection shortcut](#) [page 129]

[Data foundations with multiple connections](#) [page 135]


10.4.2 Changing a connection in a data foundation

Prerequisites

The connection you change to must be a relational connection. For multisource-enabled data foundations, the connection must also be secured.

Before you can change to a connection, you must create the local connection or connection shortcut in the local project where the data foundation is stored.

Procedure

1. Double-click the data foundation name in the local project to open the Data Foundation Editor.
2. In the Data Foundation Editor, click the **Connections** tab.
3. In the *Connections* pane, right-click the connection and select **Change...**
4. Select a new connection. How to do this depends on the data foundation type:
 - If the data foundation is single-source, the *Change Connection* dialog box lists the available connections, including the currently defined connection. Select the connection you want to change to, and click **OK**. If you have delimitation overrides set, you will be prompted for the delimitations you want to use in the new connection. For more information, see the related topic.
 - If the data foundation is multisource-enabled, the *Change Connection* dialog box displays the connection properties for the currently defined connection. Click the browse button  in the **Connection** text box. A dialog box lists the available connections. Select the connection you want to change to, and click **Finish**.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[Selecting delimitation overrides to keep](#) [page 144]

10.4.3 Selecting delimitation overrides to keep

Context

When you change the connection in a data foundation, the delimitation requirements for the new connection are automatically detected. Any delimitation overrides made for the previous connection may or may not apply. The *Select Delimitation Overrides to Keep* dialog box displays so that you can indicate which overrides to keep.

Only tables with delimitation overrides for the table name, qualifier, or owner are listed. The override value is shown in the column under **Name**, **Qualifier**, or **Owner**.

Procedure

1. For each override:
 - To keep the delimitation value set for the previous connection, leave the check box selected.
 - To remove the delimitation override and use the newly detected delimitation value, unselect the check box.

Note

A tool tip for each override gives you the following information:

- The table name, qualifier, or owner in the previous connection with the previous delimitation.
 - The new table name, qualifier, or owner in the new connection with the previous delimitation.
 - The new table name, qualifier, or owner in the new connection with the newly detected delimitation.
2. When you are finished selecting overrides to keep, click **Finish** to complete the connection change.

Results

For all other tables in the data foundation (without previous delimitation overrides) the automatically detected delimitation for the new connection applies.

Related Information


[About tables in the data foundation](#) [page 148]

10.4.4 Searching for tables in the Connection pane

Context

You can create a filtered list of tables in the **Connection** pane using a search string. You can insert tables into the data foundation from the filtered list.

Procedure

1. In the **Connection** pane of the Data Foundation Editor, click the **Show/Hide Table Search** icon, . The **Connection** pane splits. The top pane continues to display all tables in all connections. The lower pane (Search pane) displays only the tables matching the search string in the selected connection.
2. In the **Connection** pane, select the part of the connection catalog that you want to search. You can select:
 - The entire connection
 - A qualifier (if available)
 - An owner (if available)

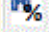
Note

For data foundations with multiple connections you can search only one connection at a time. You must select a connection. You can select the entire connection or a schema.


3. In the Search pane, enter the string you want to search for.

Note

- The wildcard character is allowed. In your search string, you must enter the wildcard character that is defined for the data source.
- Many data sources use the percent sign (%) as a wildcard, where the percent sign matches one or more characters. To enable the automatic insertion of wildcard characters (%), click the **Enable**


Automatic Wildcards icon . When automatic wildcards are enabled, if, for example, you enter the search text **2012**, the application searches for %2012%.

- The search is case-sensitive when searching connections in multisource-enabled data foundations.

4. Click the search icon  in the Search pane.
The tables with names matching the search string are listed in the Search pane. You can double-click a table in the Search pane to insert it into the data foundation.
5. To start a new search, change the search string, or select a different part of the catalog in the **Connections** pane to search, and click the search icon in the Search pane again.

Results

You can also filter the list of tables by table type using the  icon, or, in SAP HANA connections, by

information model using the . If the list of tables in the **Connection** pane is filtered, the filter applies in both the **Connection** pane and the Search pane. For more information about filtering, see the related topics.

Related Information

[About connections in the data foundation](#) [page 141]

[Filtering the tables in the connection by table type](#) [page 147]

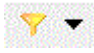
[Filtering tables in an SAP HANA connection by information model](#) [page 148]

10.4.5 Filtering the tables in the connection by table type

Context


In the **Connection** pane of the Data Foundation Editor, you can filter the list of tables in the connection by table type.

Procedure

1. Select the connection you want to filter, and open the qualifier and/or owner until tables are listed.
The application discovers table types in the connection as tables of each type are listed in the [Connection](#) pane. Once the first table type is discovered, the **Filter by Table Type** icon is available.
2. Click the down arrow next to the **Filter by Table Type** icon .
All table types discovered so far are listed. When a table type is discovered, it is automatically selected to be displayed in the connection.
3. To stop displaying a table type, unselect it in the list.
The tables are filtered in both the **Connection** pane and the table search results pane. For more information on table search, see the related topic.
4. To redisplay a table type, open the **Filter by Table Type** list and select the table type in the list.

Results

The table filter only affects the display. When you close the Data Foundation Editor, or when you change the connection, the filter selection is reset.

You can also filter the list of tables in the connection with a search string using the  icon to open the Search pane. The table type filter applies in both the **Connection** pane and the Search pane. For more information about searching for tables, see the related topic.

Related Information

[About connections in the data foundation](#) [page 141]

[Searching for tables in the Connection pane](#) [page 145]

[Filtering tables in an SAP HANA connection by information model](#) [page 148]


10.4.6 Filtering tables in an SAP HANA connection by information model

Context

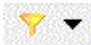
In the **Connection** pane of the Data Foundation Editor, the list of tables in an SAP HANA connection (owner `_SYS_BIC`) is filtered by default to only tables representing information models (such as an Analytic View or Calculation View). You can toggle the filter off or on.


You can also set an application preference to change the default state for the information model filter. For more information, see the related topic.

Procedure

1. Select the connection you want to filter and open the node for the `_SYS_BIC` owner.
2. To show only information models, make sure the **Filter Information Models** icon  is selected.
3. To show all tables in the `_SYS_BIC` owner, unselect the **Filter Information Models** icon.

Results

You can also filter the list of tables by table type using the  icon.

You can search the list of tables in the connection with a search string using the  icon to open the Search pane. The information model and table type filters apply in both the **Connection** pane and the Search pane.

Related Information

[Setting connection display preferences for the Data Foundation Editor](#) [page 27]

[Filtering the tables in the connection by table type](#) [page 147]

[Searching for tables in the Connection pane](#) [page 145]

10.5 About tables in the data foundation

A standard table is a graphical representation in the data foundation of a physical database table. You create standard tables when you insert database tables into the data foundation. The table and column names are inherited from the data source.

Restriction

The information design tool does not support any table or column names that are the same as an SQL reserved word. Rename these objects in the data source before inserting them into the data foundation.

Restriction

A data foundation table cannot have the same name as the qualifier or owner (see the section about qualifiers and owners). If a database table has the same name as the qualifier or owner, create an alias table with a different name.

Some data sources have tables of different types. The table type is inherited from the data source and stored as a table property in the data foundation.

Once you have inserted a standard table, you can modify it in the following ways:

- Edit table properties (name and description)
- Set the case of the table name
- Hide and unhide columns
- Change the data type of columns
- Set columns as primary and foreign keys

Tables in the data foundation can also be federated, derived, or alias tables. For more information, see the related topics.

About hidden columns

Some connections specify that certain columns in the data source are not for querying and are therefore hidden when the table is inserted into the data foundation. You can hide and unhide columns in standard tables.

Hidden columns are ignored in most workflows. For example, they do not appear in the data foundation table display, when showing table values, or when dragging and dropping the table into the business layer pane to create the related business layer objects.

If a table is joined on a hidden column, in the data foundation display, the join line points to the table header. The column is displayed when editing the join.

Hidden columns are taken into account during a refresh structure on the data foundation.

About qualifiers and owners

Some connections allow for multiple databases (called qualifiers), with different owners. Standard tables and their columns inherit the current qualifier and owner from the database. The syntax for a standard table name is as follows:

- A single-source standard table name, if the table is inserted from the current qualifier and owner, has the syntax:
"table_name"

- A single-source standard table name, if the table is inserted from a different qualifier or owner, has the syntax: `"database_qualifier"."database_owner"."table_name"`
- A multisource-enabled standard table name has the syntax:
`@catalog('short_name')."database_qualifier.database_owner"."table_name"`

i Note

Qualifier and owner are not relevant for some connections, in which case they do not appear in the data foundation connection pane and are not inherited by tables.

About delimitation

When you insert a table into the data foundation, the database requirements for delimited names is determined and the information is stored with the data foundation table properties. If the table or column names need to be delimited, the names are surrounded by double quotes in the display of the table in the data foundation.

The table, column, qualifier, and owner names that need to be delimited are surrounded by double quotes when used in an SQL expression.

In single-source data foundations, you can override the default delimitation requirements as follows:

- For standard tables you can override the delimitation on table names, column names, qualifiers, and owners.
- For alias tables you can override the delimitation on table names only. The overrides on the columns are inherited from the original table.
- For derived tables you can override the delimitation on table names only.

To override the default delimitation for tables and columns, use the **Delimit** command. For qualifiers and owners, use the **Change Qualifier/Owner** command.

When you change the connection in the data foundation, if you have delimitation overrides set, you will be prompted for the delimitations you want to use in the new connection.

Related Information

[Inserting tables into the data foundation](#) [page 151]

[Editing table properties](#) [page 152]

[Setting case of table names](#) [page 154]

[Hiding and unhiding table columns](#) [page 154]

[Changing column data types](#) [page 155]

[About table keys](#) [page 157]

[About federated tables](#) [page 198]

[About derived tables](#) [page 168]

[About alias tables](#) [page 172]

[Changing qualifiers and owners](#) [page 155]

[Changing table and column delimitation](#) [page 156]

[Selecting delimitation overrides to keep](#) [page 144]

10.5.1 Inserting tables into the data foundation


Prerequisites


Before you begin, verify that table and column names in the data source are not the same as an SQL reserved word. If so, rename these objects in the data source before inserting them in the data foundation.


Context

This procedure describes inserting tables from the data source. To insert a federated table, see the related topic.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Select **Insert Tables** from the **Insert**  menu in the data foundation view.
The *Insert Tables* dialog box lists the connections defined in the data foundation.
3. Expand the connection to see the database tables referenced in the connection.

By default, tables are listed for all qualifiers and owners. To list only tables for the currently used qualifier/owners, click the **Show Qualifiers and Owners** icon .

You can filter the list of tables in the connection by table type using the **Filter by Table Type** icon . The application discovers table types in the connection as tables of each type are listed in the pane. Once the first table type is discovered, the **Filter by Table Type** icon is available. Click the down arrow next to the **Filter by Table Type** icon to select a table type.

For SAP HANA connections, the list of tables (owner _SYS_BIC) is filtered by default to only tables representing information models (such as an Analytic View or Calculation View). Click the **Filter Information**

Models icon  to toggle the filter off and on.

4. Select a table name to insert it and all its columns into the data foundation.

Tables that are already inserted into the data foundation have an icon with a green check mark. If you insert an existing table, an alias table is inserted and you are prompted to enter a name for the alias table.

To show the values in a table, right-click the table name and select **Show Table Values**. To show the values in one column, expand the table, right-click the column name and select **Show Column Values**.

5. Select the objects you want to detect and insert automatically into the data foundation when inserting the selected tables:

Option	Description
Detect keys	Sets the key columns in the data foundation tables as they are in the database tables.
Detect row counts	Saves the number of rows in each table in the data foundation.
Detect joins	<p>Inserts the joins between the tables being inserted.</p> <p>For multisource-enabled data foundations, only joins between tables referenced by the same connection are detected. To detect joins between tables referenced by different connections, use the Detect Joins command after inserting the tables.</p>
Detect cardinalities	Saves the cardinalities of the joins as they are in the database joins.

The recommended detection options are selected by default. You can change the defaults in the application preferences. See the related topic on setting table and join detection options.

- Click **Finish** to insert the selected tables.
- Save the data foundation by clicking the **Save** icon in the main tool bar.

Results

You can also insert tables into the data foundation by dragging them from the **Connections** pane and dropping them into the data foundation view.

Related Information

- [About tables in the data foundation](#) [page 148]
- [Inserting a federated table into the data foundation](#) [page 212]
- [About alias tables](#) [page 172]
- [Detecting joins](#) [page 163]
- [Setting table and join detection options](#) [page 29]

10.5.2 Editing table properties

Context

For alias and standard tables you can edit the table name and description.

For standard tables, you can also remove columns from the table display, edit the column data types, and set or unset primary and foreign keys. The column changes you make are also made to any related alias tables.

Edit federated tables in the **Federation Layer**.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header in the data foundation view and select **Edit**.
If you are editing a derived table, the *Edit Derived Table* dialog box appears. For information on editing derived tables, see the related topic.
3. To change the table name, enter a new **Name**.

Note

When you change the name of a standard table, you break the link with the database table. For information about renaming tables using aliases, see the related link about alias tables.

4. To remove columns from the table display, unselect the columns that you want to hide and click **OK**.
This only affects the display in the data foundation view. The columns remain visible when showing table values or inserting the table into the business layer.
5. To change the data type of a column, select a data type in the list in the **Data Type** column.
The next time you refresh the data foundation structure, the original data type of the column in the database is proposed.
6. To set or unset keys, select **None**, **Primary**, or **Foreign** from the list in the Keys column.
The next time you use the **Detect Keys** command, the keys defined in the database tables override the keys set manually for a table. You can set an application preference so that if no keys are detected, the keys that you set manually in a data foundation table are kept.
7. Optionally, enter or edit the table **Description**.
8. Click **OK** to save the changes.
9. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About derived tables](#) [page 168]

[About alias tables](#) [page 172]

[Editing a federated table](#) [page 200]

[About table keys](#) [page 157]

[Setting table and join detection options](#) [page 29]

[About tables in the data foundation](#) [page 148]

[About refreshing a data foundation](#) [page 188]



10.5.3 Setting case of table names

Context

Some databases require that table names be all upper case or all lower case. Use the **Set case to** command to change the case of table names.

You cannot set the case of a federated table in the data foundation. You need to edit the table name in the **Federation Layer**.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header in the data foundation view and select  **Set case to** . Then select **Upper case** or **Lower case**.
To select multiple tables, click the table headers while holding down the **CTRL** key.
3. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[Editing a federated table](#) [page 200]

10.5.4 Hiding and unhiding table columns

Context

Hiding columns applies only to standard tables. For more information on the impact of hiding columns, see the related topic.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header in the data foundation view and select **Edit**.
3. Unselect the columns that you want to hide, select columns that you want to unhide, and click **OK**.

Note

Hidden columns are ignored in most workflows. For example, they do not appear in the display, when showing table values, or when dragging and dropping the table into the business layer pane to create the

related business layer objects. You can hide a column that is involved in a join. In this case, the column appears when editing the join.

4. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About tables in the data foundation](#) [page 148]

10.5.5 Changing column data types

Context

You can change column data types for standard tables in the data foundation.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header in the data foundation view and select **Edit**.
3. In the list of columns, select a data type in the list in the **Data Type** column.

Note

The next time you refresh the data foundation structure, the original data type of the column in the database is proposed.



4. Click **OK** to save the changes.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

10.5.6 Changing qualifiers and owners

Context

You can change the qualifiers and owners for standard tables in the data foundation. For tables in single-source data foundations, you can also change whether to delimit the qualifier and owner names.

Procedure

1. Open the data foundation in the editor by selecting it in the Local Projects View.
2. Right-click the table header in the data foundation view and select **Change Qualifier/Owner**.
To select multiple tables, click the table headers while holding down the **CTRL** key.
3. In the *Change Qualifier/Owner* dialog box, click the browse button  in the **Qualifier** field and select a new qualifier.
If the qualifier name is delimited by default, the **Delimit** option is selected. To override the default delimitation, select or unselect **Delimit**.
4. To change the owner, click the browse button  in the **Owner** field and select a new owner.
If the owner name is delimited by default, the **Delimit** option is selected. To override the default delimitation, select or unselect **Delimit**.
5. When you are finished changing the qualifier and owner information, click **OK**.
6. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information



[About tables in the data foundation](#) [page 148]

10.5.7 Changing table and column delimitation

Context

You can override the default delimitation of table and column names for tables in a single-source data foundation.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. To override the delimitation of table names and/or all column names in a table, right-click the table header in the data foundation view and select  **Delimit** .

To select multiple tables, click the table headers while holding down the **CTRL** key.

Select one of the options:

Option	Description
Yes (Tables and Columns)	Delimits the table names and all column names.
Yes (Tables only)	Delimits the table names.

Option	Description
Yes (Columns only)	Delimits all column names.
No (Tables and Columns)	Stops delimiting the table names and all column names.
No (Tables only)	Stops delimiting the table names.
No (Columns only)	Stops delimiting all column names.

Note

For alias and derived tables, you can override only table names. The columns in alias tables inherit overrides from the original table.

3. To override the delimitation of individual columns, right-click the column name and select **Delimit**.
To select multiple columns, click the columns while holding down the **CTRL** key.
4. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About tables in the data foundation](#) [page 148]

10.6 About table keys

Tables in the data foundation can have two types of keys:

Key	Description
Primary	Single or combination of columns in a table whose values identify each row in the table. The primary key guarantees row uniqueness in a table. Each table has only one primary key.
Foreign	Column or combination of columns whose values are required to match a primary or another unique key in another table. Foreign keys implement constraints, for example, not allowing a sale to be added in the Sales table for a customer that does not exist in the Customer table. Each table can have multiple foreign keys.

Keys are indicated with icons next to column in the data foundation view.

You can set keys in data foundation tables manually, or by detecting the keys in the database tables. Detecting keys does not apply to federated tables.

Related Information


[Setting and detecting table keys](#) [page 158]

10.6.1 Setting and detecting table keys

Context

You can set keys in data foundation tables manually, or by detecting the keys in the database tables. Detecting keys does not apply to federated tables.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. To set keys according to the keys detected in the database, in the data foundation view, select **Detect Keys** from the **Detect**  menu.
You can set an application preference to automatically detect keys when inserting a table into the data foundation. See the related topic on setting preferences for the Data Foundation Editor.
3. To set or unset keys manually, right-click the column in the table and select **Set as Key** and select **Primary**, **Foreign**, or **None**.
4. Save the data foundation by clicking the **Save** icon in the main tool bar.

Results

You cannot set keys for alias tables. Alias tables inherit keys from the original table.

The next time you use the **Detect Keys** command, the keys defined in the database tables override the keys set manually for a table. You can set an application preference so that if no keys are detected, the keys that you set manually in a data foundation table are kept.

Related Information

[About table keys](#) [page 157]

[Setting table and join detection options](#) [page 29]

10.7 About table row counts

Detect row counts


The number of rows in database tables can be detected and stored in the data foundation. Row counts are used to detect cardinalities in the absence of table keys.



When you detect rows counts, the number of rows for the selected tables is counted and stored.

Note

Column filters are not applied when detecting row counts.

You can also set estimated row counts for tables. This can be useful if you are working with a reduced sample of data, but want queries to be optimized for the size of the production data. The row count that you set is replaced by the detected row count when you do a detect row count for that table.

The **Detect Row Count** command on the **Detect**  menu lists the current row counts for all tables in the data foundation. From this list, you can set row counts and detect row counts for a selection of tables.

To detect row count for one table, right-click the table header in the data foundation view and select **Detect**  **Row Count** . The row count for the selected table is updated. To select multiple tables, click the table headers while holding down the **CTRL** key.

You can set an application preference to detect row counts automatically every time a table is inserted into the data foundation. See the related topic on setting table and join detection options.

Count rows

Use the **Count Rows** command on multiple tables linked by joins to see the number of rows returned by the resulting query. Column filters are applied.

To count rows returned in a query, select tables in the data foundation view in one of the following ways:

- Right-click a table and select **Select Related Tables**.
- Click the table headers while holding down the **CTRL** key.

Then right-click a table in the selection and select **Count Rows**.

Related Information

[Setting table and join detection options](#) [page 29]

10.8 About joins

A join is a condition that links tables in the data foundation. A join restricts the data returned when the two tables are queried.

Tables that are joined usually have a parent-child relationship. If tables are not joined, then a query run on the two tables can return a result set that contains all possible row combinations. Such a result set is known as a Cartesian product and is rarely useful.

Joins are defined by linking a column in one table to a column in a second table. You can insert joins into the data foundation, or detect joins automatically.

The following sections describe the types of joins you can create.

Equi-joins

An equi-join is the join type created by default between two tables. An equi-join links tables based on the equality between the values in the column of one table and the values in the column of a second table. In a normalized database, the columns used in an equi-join are often the primary key from one table and the foreign key in the other.

Self-restricting joins

A self-restricting join is when the two tables are the same. Self-restricting joins are used to define column filters. For more information on column filters, see the related topic.

Theta joins

When there is no obvious direct column to column relationship between two tables, you can use a theta join. A theta join links tables based on a relationship other than equality between two columns. It is used to link a value to a range of values. For example, an order date in one table is joined to a date between start date and end date in a second table.

Outer joins

An outer join can be used to link tables when one table contains rows that do not have a match in the common column of the other table. Unlike an equi-join, an outer join returns all rows, regardless of whether or not there is a matching value in the joined table.

A left outer join returns all rows in the first (or left-hand side) table, even if they have no match in the second table.

A right outer join returns all rows in the second (or right-hand side) table, even if they have no match in the first table.

A full outer join returns all rows from both tables, with null values when there is no match.

Shortcut joins

A shortcut join is a join that provides an alternative path between two tables. Shortcut joins improve the performance of a query by not taking into account intermediate tables, and so shortening a normally longer join path.

Shortcut joins are not taken into account to define contexts, but only to decrease the number of joins whenever possible.

Related Information

[Inserting and editing a join](#) [page 161]

[Detecting joins](#) [page 163]


[Inserting a column filter](#) [page 164]

[About contexts](#) [page 175]

10.8.1 Inserting and editing a join

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Do one of the following:

Option	Command
To edit an existing join	Right-click the join line in the data foundation view and select Edit Join .
To insert and edit a join	Select the Insert Join command from the Insert  menu in the data foundation view.

3. To define the first side of the join, select the table from the list in **Table 1**, and then select the column name. You can enter a filter pattern to filter the list of columns in Table 1. Only the column names that contain the pattern are listed.
4. To define the second side of the join, select the table from the list in **Table 2**, and then select the column name. You can enter a filter pattern to filter the list of columns in Table 2. Only the column names that contain the pattern are listed.

5. Select the join operator:

Between Table 1 and Table 2, a list of join operators lets you select how to compare the values of the columns in the join.

The default operator creates an equi-join (=). The other operators are for joins that are not based on equality between column values (>, >=, <, <=, !=).

To create a theta join using the BETWEEN operator, select the = operator. While holding down the **CTRL** key, select a second column in **Table 2**.

For more information about the types of joins possible, see the related topic about joins.

6. To create a shortcut join, select the **Shortcut join** option.

A shortcut join is a join that provides an alternative path between two tables. Shortcut joins improve the performance of a query by not taking into account intermediate tables, and so shortening a normally longer join path.

7. To create an outer join, select the **Outer join** options.

An outer join allows rows to be returned even when there is no matching row in the joined table. Select the options as follows:

To create a left outer join, select the **Outer join** option below Table 1. This join will return all rows in Table 1, even if they have no match in Table 2.

To create a right outer join, select the **Outer join** option below Table 2. This join will return all rows in Table 2, even if they have no match in Table 1.

To create a full outer join, select the **Outer join** option below both tables. This join returns all rows from both tables, with null values when there is no match.

8. Select the cardinality for the join from the **Cardinality** list. You can also click the **Detect** button to automatically detect the cardinality defined for the join in the database.

For more information about cardinality, see the related topic.

9. Optionally, edit and validate the join expression.

Based on the columns and operators you select, an SQL expression is automatically generated to define the join. You can type a custom expression for the join. To get help editing the join expression, click the **SQL**

Assistant icon .

i Note

If you change the table or column names when editing the join expression, the changes are not immediately reflected in the **Table 1** and **Table 2** lists. The changes are reflected in the lists when you save and re-edit the join.

10. Click **OK** to save the join.

11. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About joins](#) [page 160]

[About cardinality](#) [page 165]

[Inserting a column filter](#) [page 164]

10.8.2 Detecting joins

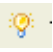
Context

Join detection looks at the data foundation tables and proposes appropriate joins. The following methods are used:

- Join detection based on column name. This method looks for identical column names in different tables. It also checks that the data type of the two columns is the same. If more than one column matches between two tables, joins are proposed for each column. Joins between a table and its alias are not proposed.
- Join detection based on database keys. This method looks for relationships defined in the database between primary keys and foreign keys.
- For data foundations with an SAP NetWeaver BW connection, join detection is based on the joins in the database schema referenced in the connection.

Before you begin, set or detect keys in the data foundation if you want to use join detection based on database keys.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Select **Detect Joins** from the **Detect**  menu in the data foundation view.
3. Select the join detection method.

For a multisource-enabled data foundation, select a method for each connection. This method is used to detect joins between tables referenced by the connection. You can also detect joins between tables from different connections. In this case, the method used is by column name.

4. From the detected joins proposed in the dialog box, select the joins to be inserted into the data foundation.
To automatically detect cardinalities for the selected joins, select the **Detect cardinalities** option.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Results

You can set an application preference to detect and insert joins automatically every time a table is inserted into the data foundation. See the related topic on setting table and join detection options.

Related Information

[Setting and detecting table keys](#) [page 158]

[About joins](#) [page 160]

[About cardinality](#) [page 165]

[Setting table and join detection options](#) [page 29]

10.8.3 Inserting a column filter

Context

A column filter, also called a self-restricting join, lets you restrict the values returned whenever the table is used in a query.

The following rules apply to column filters:

- Only one filter per column is allowed.
- You can insert a filter on a calculated column.
- The expression can contain sub-queries.
- The following @Functions are allowed in the expression: @Prompt and @Variable.
- If you insert a filter into a standard table, and then create an alias from the table, the filter is not inserted into the alias table.
- If you insert a filter into an alias tables, the filter is not automatically inserted into the original standard table.
- When you merge tables that include filters, the filters are not included in the resulting derived table.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the column name in the table in the data foundation view and select **Insert Filter**.

A self-join on the column is proposed in the [Edit Join](#) dialog box. For example, if you insert a filter on the **age** column in the **Customer** table, the following self-join is proposed:

```
"Customer"."age"="Customer"."age"
```

3. Edit the second part of the self join (the expression to the right of the equal sign) to filter the column values.

To get help editing the join expression, click the **SQL Assistant** icon .

Related Information

[About joins](#) [page 160]

10.9 About cardinality

Cardinality further describes how tables are joined by stating how many rows in one table match rows in another table. Cardinalities are needed when detecting aliases and contexts to resolve loops in the data foundation.

Cardinality of a table is expressed as a pair of numbers: the number of rows in one table that match the number of rows in the joined table. The number of rows that match can be none (0), one (1), or many (n) for each table.

For example, the tables **Customer** and **Reservations** are linked by a join.

- For each customer, there can be one or more reservations, so the cardinality of the **Customer** table is one-to-many, or 1,n.
- For each reservation, there can be one and only one customer, so the cardinality of the **Reservations** table is one-to-one, or 1,1.

Cardinality of the join is also expressed as a pair of numbers: the maximum number of rows in the second table that match one row in the first table, and the maximum number of rows in the first table that match one row in the second table.

In the example, the cardinality of the **Customer-Reservations** join is n,1, because the maximum rows that can match a row in **Customer** is n, and the maximum rows that can match a row in **Reservations** is 1.

Cardinalities can be detected for joins automatically and stored in the data foundation. The detection method first detects primary and foreign keys. Cardinalities are set according to the key status of the column in the two tables as follows:

First table column	Second table column	Cardinality
Primary key	Foreign key	1, n
Foreign key	Primary key	n,1

If no keys are detected, the cardinality is set using table row counts.

Related Information

[Detecting and setting cardinalities](#) [page 165]

[About joins](#) [page 160]


[About table keys](#) [page 157]

[About table row counts](#) [page 159]

10.9.1 Detecting and setting cardinalities

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.

2. Select **Detect Cardinalities** in the **Detect**  menu.
The *Detect Cardinalities* dialog box lists the current cardinalities for all joins in the data foundation.
3. Select the joins for which you want to detect cardinalities and click **Detect Cardinalities**.
4. To set the cardinality of a join manually, select the cardinality from the list in the **Cardinality** column for the join.
5. Click **Finish** to save the changes.

Results

You can set an application preference to detect and insert cardinality automatically every time a join is inserted into the data foundation. See the related topic on setting table and join detection options.

Related Information

[About cardinality](#) [page 165]

[Setting table and join detection options](#) [page 29]

10.10 Inserting a calculated column

Context

A calculated column is a new column in a data foundation table that is the result of a calculation based on one or more columns of the same table.

Note

Inserting a calculated time column based on a column with a time-related data type is a special case of a calculated column. For the procedure to insert a time column, see the related topic.

The following rules apply to calculated columns:

- You can insert calculated columns into standard tables only.
- You can include only columns from the same table in the SELECT statement.
- Sub-queries are not allowed.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header in the data foundation view and select **Insert Calculated Column**.
3. If the data foundation is multisource-enabled and you want to use database-specific SQL to define the calculated column, select the **Database-specific** option.

Note

Some data sources do not support database-specific SQL for defining calculated columns. In this case, the **Database-specific** option is unavailable.

For more information about SQL expressions in multisource-enabled data foundations, see the related topic.

4. Build the SQL SELECT statement that defines the column by dragging and dropping columns and functions into the **SELECT** pane.

For more information about using the SQL expression editor, see the related topic.

5. Click **Validate** to check the validity of the SQL expression.
6. Click **OK**.

The column is inserted into the table and appears in the data foundation view with a special icon. A tool tip displays the SQL expression of the calculated column when you pass the cursor over the column name.

7. To check the results of the calculated column, right-click the column and select **Show Column Values**.
8. Save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

To edit the definition, right-click the column name in the table in the data foundation view and select **Edit Calculated Column**.

Related Information

[Inserting a time column](#) [page 168]

[About the SQL/MDX Expression Editor](#) [page 358]

[SQL expressions in multisource-enabled data foundations](#) [page 136]

10.11 Inserting a time column

Context

A time column is a calculated column that contains a date part (for example, month, quarter, or year) based on a column with a time-related data type.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click a column with a time-related data type and select **Insert Time Column**.
Columns with time-related data types have a special icon that resembles a calendar.
3. Select a date part from the list.
A calculated column is inserted into the table and appears in the data foundation view with a special icon. A tool tip displays the SQL expression of the calculated column when you pass the cursor over the column name.
4. To check the results of the calculated column, right-click the column and select **Show Column Values**.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

To edit the definition, right-click the column name in the table in the data foundation view and select **Edit Calculated Column**.

Related Information

[Inserting a calculated column](#) [page 166]

10.12 About derived tables

A derived table is a virtual table in the data foundation that combines other tables using calculations and functions. You can create objects in the business layer on a derived table in the same way that you do for a standard table. Use derived tables in the following situations:

- To create a table with columns from other tables. The column definitions can include complex calculations and functions.

- To create a single table that combines two or more tables (called merging tables).
- To create a table that contains a selection of columns from different tables.

Related Information

[Inserting a derived table based on a data foundation table](#) [page 169]

[Merging tables](#) [page 170]

[Inserting and editing a derived table](#) [page 171]

10.12.1 Inserting a derived table based on a data foundation table

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header of the table to be the basis for the derived table, and select **Insert > Derived Table**.
3. Enter a name for the derived table that is unique in the data foundation, and click **OK**.

Results

A derived table with the new name and all the columns from the original table is inserted into the data foundation.

Next Steps

Edit the derived table to make the desired modifications.

Related Information

[Inserting and editing a derived table](#) [page 171]

[About derived tables](#) [page 168]

10.12.2 Merging tables

Context

Merging tables inserts a derived table into the data foundation that consists of the combined columns from two or more tables linked by joins. Federated tables cannot be merged.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. In the data foundation view select the tables you want to merge:

Option	Command
To select a table and all tables related to it by joins	Right-click the table and select Select Related Tables .
To select tables manually	Click the table headers while holding down the CTRL key.

3. Right-click the selection of tables and select **Merge**.
4. Enter a name for the table that is unique within the data foundation and click **OK**.
The merged table is inserted as a derived table. The new table is joined to any tables that the original tables were joined to.
5. Select if you want to delete the original tables.
The original tables become obsolete and you have the choice of deleting them. If you choose to keep the original tables, the joins linking those tables are deleted, however the tables stay in the data foundation.

Results

In a multisource-enabled data foundation, a derived table that is the result of a merge creates expressions using SQL-92 standard syntax. To use database-specific SQL, you must edit the derived table and explicitly select database-specific syntax.

Next Steps

To edit the merged table, right-click the table header and select **Edit**.

Related Information


[Inserting and editing a derived table](#) [page 171]

[About derived tables](#) [page 168]

10.12.3 Inserting and editing a derived table

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Do one of the following:

Option	Command
To edit an existing derived table	Right-click the table header in the data foundation view and select Edit .
To insert and edit a derived table	Select the Insert Derived Table command from the Insert  menu in the data foundation view.

3. Give the derived table a unique name within the data foundation.
4. In a multisource-enabled data foundation, if you want to include database-specific functions in the definition of the derived table, select the **Database-specific** option.

For more information about SQL syntax in multisource-enabled data foundations, see the related topic.

5. Enter or edit the SQL expression for the derived table in **Expression**.

Note

To build an expression for the first time, you can use the **SQL Builder**. The SQL Builder works like the query panel. You drag and drop tables and columns to be included in the derived table. The SQL expression is generated automatically.

For more information on using the SQL Expression Editor, see the related topic.

6. Click **Validate** to check the validity of the SQL expression.
7. Click **OK**.
8. Link the derived table to other tables in the data foundation by inserting the appropriate joins.
9. To check the results of the derived table, right-click the table and select **Show Table Values**.
10. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[SQL expressions in multisource-enabled data foundations](#) [page 136]

[About the SQL/MDX Expression Editor](#) [page 358]

[Inserting and editing a join](#) [page 161]

[Using analytical functions in a derived table definition](#) [page 240]

To use analytic functions in the data foundation, you define the analytic function in the SELECT statement for a derived table.

10.13 About alias tables

An alias table is a reference to a standard, derived, or federated table in the data foundation. It is an identical duplicate of the original table (except for column filters), but has a different name. The data in the table is exactly the same as the original table, but the different name "tricks" the SQL of a query to accept that you are using two different tables.

You use alias tables to break loops in the join paths in the data foundation. The **Detect Aliases** command analyzes the join paths and proposes alias tables to break any loops detected in the data foundation. For more information about resolving loops, see the related topic.

You also use alias tables to rename a table. The link between the data foundation and database is based on the table name. If you create an alias to give the table a new name, the link to the database table is preserved, but the alias table name is used in the data foundation.

To find alias tables already inserted into the data foundation, you can do a data foundation search. There are also commands to highlight the alias tables of an original table, and highlight the original for an alias table. For more information, see the related topics.

Related Information

[Detecting alias tables](#) [page 173]

[Inserting alias tables](#) [page 172]

[Resolving loops](#) [page 177]

[Searching for tables and columns in the data foundation](#) [page 191]

[Highlighting aliases](#) [page 174]

[Highlighting the original table of an alias](#) [page 174]

10.13.1 Inserting alias tables

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. In the data foundation view, select the table to be the basis of the alias.
You can create aliases for more than one table at a time. Click the table headers while holding down the **CTRL** key.
3. Right-click the selection and select **Insert > Alias Table**.
4. In the [Insert Alias Tables](#) dialog box, unselect any aliases you do not want to insert.
5. Edit the names for the alias tables in the **Alias Name** column, and click **OK**.

Results

The selected alias tables are inserted into the data foundation. The original table name is listed in parentheses in the table header.

Next Steps


To edit the name and description of an alias table, click the table header in the data foundation view and select **Edit**.

Related Information

[About alias tables](#) [page 172]

10.13.2 Detecting alias tables

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Select **Detect Aliases** from the **Detect**  menu in the data foundation view.

You can also detect aliases from the **Aliases and Contexts** pane in the Data Foundation Editor. Click the

Detect Aliases icon .

The command analyzes the join paths and proposes alias tables to break any loops detected in the data foundation.

3. If alias tables are proposed, select which aliases you want to insert automatically.

Results

The selected alias tables are inserted into the data foundation. The original table name is listed in parentheses in the table header.

Next Steps

To edit the name and description of an alias table, click the table header in the data foundation view and select **Edit**.

Related Information

[About alias tables](#) [page 172]

10.13.3 Highlighting aliases

Context


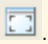
Use this command to highlight the alias tables associated with a standard or derived table in the data foundation.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header of the original table and select **Highlight Aliases**.

The original and any associated alias tables are highlighted. All other tables are grayed.

Note

Some alias tables may be off the visible area of the data foundation view. You can quickly check for hidden highlighted tables by clicking the **Fit to Window** icon at the bottom of the data foundation view . To undo the **Fit to Window**, click the **Reset Zoom** icon .

3. To return to the normal data foundation view display, click anywhere in the view.

10.13.4 Highlighting the original table of an alias

Context



Use this command to highlight the original table in the data foundation of an alias table.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header of the alias table and select **Highlight Original Table**.

The original and alias tables are highlighted. All other tables are grayed.

Note

The original table may be off the visible area of the data foundation view. You can quickly check for a hidden table by clicking the **Fit to Window** icon at the bottom of the data foundation view . To undo the **Fit to Window**, click the **Reset Zoom** icon .

3. To return to the normal data foundation view display, click anywhere in the view.

10.14 About contexts

A context is a collection of joins which provide a valid query path. The most common use of contexts is to resolve loops in the data foundation when the loop cannot be resolved by creating an alias table. Another use of contexts is when multiple fact tables share a dimension table. In this case, a context is created for each fact table.

In the information design tool, a context resolves a loop by identifying a set of joins that define one specific join path through the tables in the loop. The user is prompted for the context to use at query time. The context ensures that joins are not included from different paths within the same SQL query.

A context is defined by setting states for the joins involved in the ambiguity. In a context, a join has one of three states:

- Included joins: In a part of the schema that is ambiguous, the context solves the loop by defining a path with the included joins.
- Excluded joins: In a part of the schema that is ambiguous, the excluded joins define the path that context will never take.
- Neutral joins are in a part of the schema that is not ambiguous, and are always included in the query path of the context. Any join that is not explicitly included or excluded is neutral.

When a new join or table is inserted into the data foundation, it is neutral by default. Contexts will not need to be updated unless the new table or join is explicitly involved. You can change the default so that added joins are automatically excluded or included. You change this default behavior in the application preferences for the Data Foundation Editor. You can also choose to use the new default behavior when adding contexts.

You can insert contexts into the data foundation manually, or by detecting contexts. The detect command analyzes the join paths and proposes contexts to resolve any loops that cannot be resolved by alias tables.

Related Information

[Detecting contexts](#) [page 176]

[Inserting and editing contexts](#) [page 177]

[Resolving loops](#) [page 177]


[Setting default join states for contexts](#) [page 31]

10.14.1 Detecting contexts

Prerequisites

Before detecting contexts, you need to set cardinalities and detect aliases. See the related topic about resolving loops for the prerequisite tasks.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. From the **Aliases and Contexts** pane, select the **Detect Contexts** icon .

The command analyzes the join paths and proposes contexts to resolve any loops that cannot be resolved by alias tables.

Note

You may get a message that the loop can be resolved using aliases. See the related topic on resolving loops.

3. In the *Detect Contexts* dialog box, select the contexts you want to insert.

To see the context highlighted in the data foundation view, click the proposed context name. A join that is

included in the context is shown by an included icon . A join that is excluded is shown by the excluded

icon .

4. Click **OK** to insert the selected contexts into the data foundation.
The new contexts are listed in the **Aliases and Contexts** pane in the **Contexts** folder.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information


[About contexts](#) [page 175]

[Resolving loops](#) [page 177]

10.14.2 Inserting and editing contexts

Procedure



1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Select the **Aliases and Contexts** pane.
3. Do one of the following:

Option	Command
To edit an existing context	Select the context in the Contexts folder.
To insert and edit an context	Select the Insert Context icon  .

The properties of the context are displayed in the **Context Properties** pane:

- Context name
 - All joins in the data foundation
 - The state of the join in this context: whether the join is included, excluded, or neutral
4. Edit the context name in **Name**.
 5. To include or exclude a join, or to set the join to neutral, click the join expression in **Join Expression** list. The state toggles each time you click.

You can also toggle the state by clicking the join line in the data foundation view.

A join that is included in the context is shown by an included icon . A join that is excluded is shown by the excluded icon .

6. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About contexts](#) [page 175]

10.15 Resolving loops

Context


Loops occur when multiple paths join tables. The rows that are returned from the query are the intersection of the results for each path, so fewer rows are returned than expected.


An alias table breaks a loop by using the same table twice in the query, once for each path. That way, the rows returned in the query are the union of the results for each path.

A cycle is a loop that occurs when the tables joined by the loop all have a cardinality of (1,n). In this case, the **Detect Aliases** command cannot determine which table to create an alias for.

When loops cannot be resolved with an alias table, contexts are used. You use contexts to resolve the ambiguity by explicitly directing the query as to which join path to use.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Detect and insert all the joins in the data foundation.
3. Detect or set the cardinalities for the joins.
4. Make sure there are no cycles in the data foundation. You can do this by examining the cardinalities of any loops, or by refreshing the loop resolution status (see the following steps in this procedure).
5. Make sure none of the joins have a cardinality of (n,n). Aliases and contexts cannot be detected. Set the cardinality manually for (n,n) joins.
6. From the **Aliases and Contexts** pane, detect aliases.
7. Detect contexts. It is recommended to insert all proposed contexts.
8. In the **Loops** box, click the **Visualize Loops** icon  to check if all the loops have been resolved.

Possible loops are listed in **Loops**. To check if the loops are resolved, click the **Refresh Loop Resolution Status** icon .

A message appears suggesting what to do about unresolved loops.

A loop is resolved when you see a green check mark next to the loop name.

Related Information

[Detecting joins](#) [page 163]

[Detecting and setting cardinalities](#) [page 165]

[Detecting alias tables](#) [page 173]

[Detecting contexts](#) [page 176]

10.16 About input columns in the data foundation

An input column is a parameter in the data source that expects a value. The parameter is represented in a table column in the data foundation.

For each input column, you can provide a static value, or a parameter defined in the data foundation. The parameter can prompt the user for a value and be associated with a list of values. For some input columns, providing a value is optional.

Following are examples of input columns in the data foundation:

- SAP NetWeaver BW key date variables. An input column is inserted into each table in the data foundation that handles time-dependent data. To resolve these input columns at query time, a parameter is inserted into the data foundation called key date. Because SAP NetWeaver BW key date variables are mandatory, by default, at query time, the key date parameter is not prompted. It is automatically assigned the current date. You can edit the prompt parameters in the data foundation.
- SAP ERP ABAP function input parameters. One data foundation table is created to map the main function. It contains input columns for the input parameters of the function. These parameters can be mandatory or optional. For mandatory parameters, you enter either a static value or a data foundation parameter for the associated input column.

Related Information

[Editing input columns](#) [page 179]

[Using SAP NetWeaver BW data sources](#) [page 40]

10.16.1 Editing input columns

Prerequisites

To assign a parameter to an input column, you must first define the parameter in the data foundation. For SAP NetWeaver BW key date variables, a data foundation parameter is automatically inserted. For more information on parameters, see the related topic.


Procedure

1. You can list input columns to edit in three ways:

Option	Command
To list input columns for a table	Right-click the table header in the data foundation view and select Edit Input Columns .
To list input columns for all tables	Right-click anywhere in the data foundation view and select Edit Input Columns .
To edit a single input column	Right-click the column name in the data foundation view and select Edit Input Column .

If no input columns exist in the table or data foundation, the **Edit Input Column** command is not available.

2. To assign a value to an input column, in the [Edit Input Columns](#) dialog box, select the column in the list.
 - **No assignment** is selected by default and means that no value or parameter is assigned to the column.
 - To assign a static value, select **Value** and enter a value in the text box. You can enter blanks or leave the text box empty to assign a blank value to a column with character data type.

- To assign a parameter, select **Parameter**. Click the icon  to select from a list of parameters defined in the data foundation.

Note

The **Values** column indicates mandatory input columns as **[mandatory]**, and optional input columns as **[optional]**.

3. To assign the values you entered, click **OK**.
4. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About parameters](#) [page 280]

10.17 About parameters and lists of values in the data foundation

A parameter is a variable in the universe that requires a value at query time. Parameters are often defined to prompt the user to supply a value, and in this case are referred to as prompts.

A list of values is a collection of data values that can be associated with an object in the universe, allowing the user to choose values for a prompt.

You can insert parameters and lists of values into the data foundation. They are inherited by any business layer built on the data foundation, but cannot be modified in the business layer.

To insert a parameter or list of values, go to the **Parameters and Lists of Values** tab in the Data Foundation Editor. From there, the procedure is the same as inserting parameters and lists of values into a business layer. See the related topics.

Related Information

[About parameters](#) [page 280]

[About lists of values](#) [page 283]

10.18 About data foundation properties

The following properties apply to the entire data foundation:

Property	Description
Description	Describes the data foundation. The description can be entered when you create the data foundation in the New Data Foundation wizard, and edited at any time in the data foundation properties.
Allow Cartesian products	<p>When selected, if the SQL expression that defines an object in the data foundation could result in a Cartesian product, the SQL is allowed.</p> <div> i Note A Cartesian product is a result set which contains all the possible combinations of each row in each table included in a query. A Cartesian product is almost always an incorrect result. </div>
Multiple SQL statements for each context	When selected, allows the user to select the query path when the query involves contexts. This option should be selected if the data foundation contains contexts.
SQL Parameters	Specifies custom values for SQL generation parameters that override the default values.
Comments	Contains comments about the data foundation.
Summary	Displays a summary of the number of each type of object defined in the data foundation.

Related Information

[Editing SQL options in the data foundation](#) [page 181]

[Setting SQL generation parameters in the data foundation](#) [page 182]

[Showing a data foundation summary](#) [page 183]

[About contexts](#) [page 175]

10.18.1 Editing SQL options in the data foundation

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation in the Local Projects View.
2. Make sure the top level of the data foundation is selected in the tree view in the **Data Foundation** pane.
3. Click the **SQL Options** tab in the properties pane.
4. Select or unselect options as required. For a description of the options, see the related topic.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About data foundation properties](#) [page 180]

10.18.2 Setting SQL generation parameters in the data foundation

Context

The custom values for SQL generation parameters in the data foundation override the default values.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation in the Local Projects View.
2. Make sure the top level of the data foundation is selected in the tree view in the **Data Foundation** pane.
3. Make sure the **Properties** tab is selected in the properties pane.
4. Click the **Parameters** button.
5. In the [Query Script Parameters](#) dialog box, edit the parameters:

The currently defined SQL generation parameters are listed. Non-default parameters and parameters with non-default values are in bold type.

Option	Command
To change the value of an existing parameter	Click the Value column and select or enter the new value.
To add a predefined parameter	Click the arrow in the list box next to the Add button to display the list of predefined parameters. Select the parameter from the list and click Add .
To add a custom parameter	Make sure no predefined parameter is listed in the box next to the Add button, then click Add . A parameter with a default name is added to the table. To edit the parameter name, click the Name column. Click the Value column to enter a value.

To see a description of all the predefined SQL generation parameters and their values, click the help button.

6. To return to the default list of parameters and the default values, click **Default Values**. This removes any added parameters from the list and sets all values to the default.

Related Information

[About SQL Generation Parameters](#) [page 433]

10.18.3 Showing a data foundation summary

Context

Use this command to show the number of each type of object defined in the data foundation.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation in the Local Projects View.
2. Make sure the top level of the data foundation is selected in the tree view in the **Data Foundation** pane.
3. Make sure the **Properties** tab is selected in the properties pane.
4. Click the **Summary** button.

Results

The data foundation summary displays in a new dialog box.

10.18.4 Editing the data foundation description and comments

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Make sure the top level of the data foundation is selected in the tree view in the **Data Foundation** pane.
3. To enter or edit a description, click the **Properties** tab in the properties pane.
4. To enter or edit comments, click the **Comments** tab in the properties pane.

The comments apply to the entire data foundation. You can also enter comments in the data foundation display. For more information, see the related link.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About data foundation properties](#) [page 180]

[Inserting a comment into the data foundation view](#) [page 192]

10.18.5 Showing SAP HANA variable information

Context

Variables and input parameters in SAP HANA information models are automatically associated with the corresponding tables in the data foundation. You can see variable information in the properties pane.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation in the Local Projects View.
2. Make sure the **Data Foundation** pane is selected.
3. Select the table in the data foundation that corresponds to the information model for which you want to see the variables.
4. Select the **Variables** tab in the properties pane.

Related Information

[Using SAP HANA data sources](#) [page 45]

10.19 Showing table values

Context

You can show values for one or more tables in a data foundation. If filters are defined on any of the columns, the filters are applied when showing values. To see values in the database for a table (no data foundation filters are applied), show the values on a table in the Connections pane.

The show values command by default opens a tab in the editor to display the values. You can set a preference to have the values open in a dedicated view or a dialog box. For more information, see the related topic.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
You can also show table values from the data foundation view in the Business Layer Editor. In this case, open the business layer.
2. Do one of the following:

Option	Command
To show values for one or more tables with filters applied	Right-click the table header in the data foundation view. To select multiple tables, click while holding down the CTRL key.
To show values for one table without filters applied	From the Connections pane in the Data Foundation Editor, expand the connection and right-click the table name.

3. Select **Show Table Values**.

The show values window appears. To see what you can do in this window, see the related topic on showing values in a data source.

Related Information

[Showing values in a data source](#) [page 185]

[Setting preferences for showing values](#) [page 34]

10.19.1 Showing values in a data source

You can show values in the underlying data source for data foundation tables and columns, business layer objects, and connections. This topic describes what you can do when showing values.

Note

Unless you are showing values from the connection, any column filters defined in the data foundation are applied when retrieving values.

To limit the number of rows returned from the data source, enter a number in **Max Rows**.

To see the query script, click **View Log**.

What you can do when showing values in the **Raw Data** tab:

- Re-order the columns in the display: Drag and drop column headers to the new location in the table.
- Sort rows by column: Click the column header to sort the rows by the column value in ascending or descending order.
- Filter rows by column: Click **Add Filter** and construct a filter for one or multiple columns using the filter value selector.
- Filter results to rows that contain a character or group of characters in any of the columns: Enter the characters to filter on in the **Enter your filter** text box. You can use the * character as a wildcard, for example:
 - If you enter B in the filter text box, only rows with a column value that contains the character B display.
 - If you enter B*, only rows with a column that contains a value starting with B display.
 - If you enter *B, only rows with a column that contains a value ending with B display.
- Export the results to a local file (.csv or .xml format): Click **Save as File**.

To see distinct values for a selected column, click the **Distinct Values** tab and select a column.

To build and format a chart, click the **Analysis** tab. To save the chart as an image, click **Save As Image**.

Related Information

[Showing table values](#) [page 184]

[Showing column values](#) [page 186]

[Showing business layer object values](#) [page 277]

10.20 Showing column values

Context

You can show values for one or more columns in a data foundation table. If a filter is defined on the column, the filter is applied when showing values. To see values in the database for a column (no data foundation filters are applied), show the values on a column in the Connections pane.

The show values command by default opens a tab in the editor to display the values. You can set a preference to have the values open in a dedicated view or a dialog box. For more information, see the related topic.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View. You can also show column values from the data foundation view in the Business Layer Editor. In this case, open the business layer.

2. Do one of the following:

Option	Command
To show values for one or more columns with filters applied	Right-click the column name in the data foundation view. To select multiple columns, click while holding down the CTRL key.
To show values for one column without filters applied	From the Connections pane in the Data Foundation Editor, expand the connection and right-click the column name.

3. Select **Show Column Values**.

The show values window appears. To see what you can do in this window, see the related topic on showing and profiling values in a data source.

Related Information

[Showing values in a data source](#) [page 185]

[Setting preferences for showing values](#) [page 34]

10.21 Profiling column values

Context

You can profile the values for a column in a data foundation table. Profiling shows graphically (in a pie or bar chart) the number of occurrences of each value of a column. If the column has a filter defined, the filter is applied.

Procedure

1. You can profile column values from the data foundation view in either the data foundation or business layer editor. Open the editor by double-clicking the resource in the Local Projects View.
2. In the data foundation view, right-click the column name in the table display and select **Profile Column Values**.
The profiled data is shown in a table.
3. To see the profiled data as a chart, select the **Pie Chart** or **Bar Chart** option.

10.22 Showing local dependencies in the data foundation

Context

Use the **Show Local Dependencies** command when you plan to change tables and columns in the data foundation. The command will find the business layers and their objects that depend on the table or column.

Procedure


1. Right-click the table header or column name in the data foundation view and select **Show Local Dependencies**.
You can select several tables and/or columns by holding down the **CTRL** key.
The business layers that depend on the selected tables and columns are listed.
2. Select the business layer for which you want to see the dependent objects.
A dialog box lists the data foundation tables and columns, and the business layer objects that depend on them.
3. To edit a business object, double-click the object name in the **Business layers and objects** box. The Business Layer opens with the focus on the selected object.

Related Information

[About resource dependencies](#) [page 318]

10.23 About refreshing a data foundation

Refreshing the structure of a data foundation compares the existing tables in the data foundation with those in data source and proposes updates to the data foundation tables: deletes obsolete tables and columns, inserts missing columns, and updates changed columns.

To start the Refresh Structure Wizard, in the Data Foundation Editor, select **Refresh Structure** from the **Detect**  menu.

The wizard detects the following changes and lists them each in their own dialog box. In each case, you select which of the proposed changes to make in the data foundation.

- Tables in the data foundation that were deleted in the database. The wizard proposes to delete these tables and any related joins from the data foundation.
- Columns in data foundation tables that were deleted in the database tables. The wizard proposes to update each corresponding table in the data foundation to delete these columns and the joins that use these columns.
- Columns added in the database. The wizard proposes to update each corresponding table in the data foundation to add these columns.
- Column data types changed in the database. The wizard proposes to update the data type of each column in the data foundation that is different from the database column type.
- For connections to SAP HANA, variables in the data source that are added, deleted, or modified.

The wizard lists your selected changes in a summary dialog box and asks for confirmation before continuing with the refresh.

After refreshing the structure, save the data foundation by clicking the **Save** icon in the main tool bar.

Note

For data foundations based on SAP NetWeaver BW connections, you can detect new tables and joins in the data source and insert them into the data foundation using the **Synchronize Tables** command.

Related Information

[Synchronizing tables](#) [page 189]

10.23.1 Synchronizing tables

Prerequisites

Before synchronizing tables, refresh the data foundation structure to make sure all existing data foundation tables are updated with any new columns in the data source.

Context

Synchronizing tables applies only to multisource-enabled data foundations based on SAP NetWeaver BW data sources.

Synchronizing tables searches the data source for new tables (using the SAP NetWeaver BW strategy), and inserts the new tables and joins into the data foundation.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Click the [Connections](#) pane.
3. Right-click the connection in the [Connections](#) pane and select **Synchronize Tables**.
4. You are prompted to optionally detect new joins.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

The insertion can be undone using the Edit > Undo command on the main menu.

Next Steps

Refresh the business layer with new objects in the data source using the **Insert Candidate Objects** command.

Related Information

[About refreshing a data foundation](#) [page 188]

[Inserting candidate objects](#) [page 293]

[Refreshing universes based on SAP Netweaver BW](#) [page 44]

10.24 Inserting a custom data foundation view




Context

A custom data foundation view is a subset of the data foundation **Master** view. Use views when you are editing a data foundation that contains many tables, and you are interested in working with a subset of tables. You can define multiple custom views for the data foundation.

Table operations are allowed from all views. Any changes to a table, for example assigning the table to a family, are propagated to all views in the data foundation.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Do one of the following:

Option	Command
To insert an empty view	Select Insert View from the Insert  menu.
To insert a view based on a selection of tables	Select one or several tables (click the table headers while holding down the CTRL key). Right-click the selection and select ► Insert ► View from Selection  .
To insert a view based on search results	Use the Search Panel to search for the tables to include in the view. For more information on how to search, see the related topic. Right-click the selection in the search results view and select ► Insert ► View from Selection  .




3. Enter a name for the view and click **OK**.
A new tab appears at the bottom of the view pane and the new view is displayed.
4. To add a table to a view:
 - a) Click the **Master** view tab, or any other view that contains the table you want to add.
 - b) Select the table or tables you want to add.
 - c) Right-click the table header and select **Add to View**.
 - d) Select the view from the list (only views that do not already contain the table are listed).
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[Searching for tables and columns in the data foundation](#) [page 191]

10.25 Searching for tables and columns in the data foundation

Procedure

1. To open the search panel, in the data foundation view, click the **Show/Hide Search Panel** icon .
2. By default, the search looks for tables. To look for columns, click the icon  in the filter text box.
3. You can limit your search in several ways:
 - Enter text to search for in the filter text box.
 - Select connections, table types, column types, families, and contexts in the respective lists.The tables that match the search criteria are highlighted in the data foundation view.
4. To modify the view to display only the matching tables, click the **Search Options** icon  at the top of the search panel, and select **Auto arrange search results**.

→ Tip

You can also use the **Center on Selection** command to change the zoom on the data foundation display so that all the tables in a selection are visible in the display window.

5. Click **Reset** to clear the search criteria and start a new search.

Results

Certain operations on tables are not possible when the Search Panel is active, for example inserting alias and derived tables, detecting joins, or checking integrity. The data foundation commands that are not available when using the Search Panel are grayed. To use these commands, close the Search Panel by clicking the **Show/Hide**

Search Panel icon .

Related Information


[Centering the view on a selection](#) [page 192]

10.26 Inserting a comment into the data foundation view

Context

A comment is a note that you can place anywhere in a data foundation view.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation in the Local Projects View.
2. Select **Insert Comment** from the **Insert**  menu.
3. In the *Edit Comment* dialog box, define the display parameters of the note, and enter the comment text.
The comment is inserted into top left corner of current view.
4. Drag the comment to the location in the view where you want it to appear.
5. Save the data foundation by clicking the **Save** icon in the main tool bar.

10.27 Centering the view on a selection


Context

The **Center on Selection** command lets you temporarily change the zoom on the data foundation view display so that all the tables in a selection are visible in the display window.

Procedure

1. In the data foundation editor, make a selection of tables.
For example, using the Search Panel, select all tables in a particular family.
2. Right-click the table header of one of the selected tables and select **Center on Selection**.

Results

The data foundation display zooms so that all the selected tables appear in the display window. To reset the display, close the Search Panel if it is open, or click the **Reset Zoom** icon  in the lower toolbar of the data foundation view.

Next Steps

➔ Tip

You can also center the display on a table or join by selecting the table or join name in the tree view in the **Data Foundation** panel to the left of the display.

Related Information

[Searching for tables and columns in the data foundation](#) [page 191]

10.28 Changing the display of objects in the data foundation

The following commands let you change the display of objects (tables, columns, and joins) in the data foundation view. You can also set application preferences that affect the display of data foundation objects. For more information on any topic, click the link.

- [Auto-arranging tables in the data foundation view](#) [page 193]
- [Changing table display](#) [page 194]
- [Grouping tables using families](#) [page 194]
- [Setting display preferences for the data foundation view](#) [page 28]
- [Setting performance-related options for the data foundation view](#) [page 31]


10.28.1 Auto-arranging tables in the data foundation view

Prerequisites

Once you have inserted tables and joins into the data foundation view, you can automatically arrange the tables according to the flow of joins from one to many.

To change the display of individual tables, see the related links.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Click the **Auto Arrange Tables** icon  in the data foundation view.

The tables are arranged in the view.

3. Save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[Changing table display](#) [page 194]

[Hiding and unhiding table columns](#) [page 154]

[Grouping tables using families](#) [page 194]

10.28.2 Changing table display

Context

For each table in the data foundation, you can select how much table information to display in the view.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Right-click the table header in the data foundation view and select **Display**.
3. Select a display mode:

Option	Description
Collapsed	Displays only the table header.
Joins only	Displays the table header and any columns that are part of a join.
Expanded	Displays the table header and all columns.

You can also toggle through the different display modes by clicking the arrow icon on the right side of the table header.

4. Save the data foundation by clicking the **Save** icon in the main tool bar.


10.28.3 Grouping tables using families

Context

A family is a set of display parameters that can be used to visually group tables of the same type. For example, you might define different families for fact and dimension tables.

The display parameters include background color, text color and font.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Click the **Edit Families** icon  in the data foundation view.
3. In the *Edit Families* dialog box, create a family for each table group:
 - a) Click **Add** to define a new family.
 - b) Enter a name for the family in **Name**.
 - c) Edit the table color, background, and font for the family.
 - d) Click **Apply** to save the family definition.

You can export and import family definitions. Export creates a file in a local folder that can be shared between different users of the information design tool.

4. When you have added all families, click **OK**.
5. Assign tables to families. For each family:
 - a) Select the tables to be assigned to one family. Click the table headers while holding down the **CTRL** key.
 - b) In the **Families** list in the data foundation view tool bar, select the family.

Note

When you assign a table to a family, the table acquires the display attributes of the family in the current view and all data foundation views where the table is present.

6. To remove a table from a family, select the table, and in the **Families** list, select **No family**.
7. Save the data foundation by clicking the **Save** icon in the main tool bar.

11 Working with the federation layer

11.1 About the federation layer

The federation layer is available only in multisource-enabled data foundations. It lets you create federated tables that can include data from any of the data source connections defined in the data foundation. Federated tables can be inserted into the data foundation and used to define the schema on which the universe is built.

At design time, you use the federation layer to define a data flow composed of data source tables and federated tables. You define the data flow graphically without needing to write a lot of detailed SQL statements. You can specify complex transformations of data within the flow, and build multi-level data flows by using a federated table as input to another federated table.

The federation layer lets you maintain a coherent set of federated tables. From this set, you selectively insert federated tables into the data foundation.

Related Information

[Building the federation data flow](#) [page 196]

11.2 Building the federation data flow

Prerequisites

Before you can build the federation data flow, you must have a multisource-enabled data foundation based on at least one valid connection.

Context

Building the federation data flow is a matter of designing the input flow into a coherent set of federated tables that your applications will query.

In a top-down approach to the design, you start with the final schema of federated tables. You define these federated tables, then define the inputs and mapping for each column.

In a bottom-up approach, you start with the data source tables. You add federated tables from a data source table, and then modify the mappings.

The following procedure describes the steps to build a federation data flow. Links to more information on each step in the procedure can be found in the Related Topics.

Procedure

1. You build the federation data flow using the Data Foundation Editor. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. To open the **Federation Data Flow** view, click **Federation Layer**.
3. Add a federated table either manually, or from a data source table.
4. Define the input to the federated table. An input table can be a data source table, or a federated table.
 - If you added a federated table manually, add one or more input tables and join them.
 - If you added a federated table from a data source, you can add other input tables and join them.
5. Map the columns from the input tables to the federated table.
6. You can further refine the mapping by editing the mapping formulas, adding pre-filters and post-filters, and specifying distinct rows for input tables.
7. Optionally, you can define additional mappings for the federated table.

Mappings can be activated and deactivated. When more than one mapping is activated, the effective mapping is a union of all activated mappings.
8. Repeat the steps to add other federated tables to your data flow.
9. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

Check the integrity of the federation layer often during the design of federated tables. After you have created and validated a federated table, you can insert it into the data foundation.

Related Information

[About the Data Foundation Editor](#) [page 136]

[Adding a federated table manually](#) [page 199]

[Adding a federated table from a data source](#) [page 200]

[About input tables and joins](#) [page 201]

[Mapping columns from the input table to columns of the federated table](#) [page 206]

[Editing a mapping formula](#) [page 206]

[About pre-filters and post-filters](#) [page 209]

[Adding a mapping](#) [page 207]

[Activating and deactivating mappings](#) [page 208]

[Checking integrity of the federation layer](#) [page 212]

[Inserting a federated table into the data foundation](#) [page 212]

[About distinct rows on input tables](#) [page 209]

11.3 About federated tables

Federated tables are the tables that you create to present data in the correct format for your data foundation. A federated table may be the end result, or a table that contributes to a federated table at a higher level.

A federated table can be added in two ways:

- A federated table added manually is empty. You add the columns and define their properties.
- A federated table added from a data source by default contains the same columns as the data source table. The columns inherit the properties from the data source.

You edit a federated table to add and delete columns and change the column properties.

Federated table columns have the following properties:

Property	Description
Name	The default column name can be edited.
Data Type	The data type of the column can be selected from a list.
Input	<p>Whether or not the column expects input. The input can be optional or mandatory.</p> <p>Input columns can be resolved in the federation layer by a join or filter. Input columns that are not resolved in the federation layer, are resolved in the data foundation.</p> <p>When showing table values for a federated table in the data foundation or federation layer, you are prompted to enter values for the input columns. A message indicates if a value is mandatory or optional. Select the input column and enter a value in Assignment.</p>
Description	An optional description for the column.

You can also describe additional logic by adding distinct rows on input tables, pre-filters, joins between input tables and post-filters.

The logic you build into a mapping is applied in the following order:

1. Distinct rows
2. Pre-filters
3. Input table joins
4. Post-filters
5. Mapping formulas

A federated table can have more than one mapping. All mappings are activated by default. When more than one mapping is activated, the effective mapping is a union of all activated mappings.

Related Information

[About input columns in the data foundation](#) [page 178]

[Adding a federated table manually](#) [page 199]
[Adding a federated table from a data source](#) [page 200]
[Editing a federated table](#) [page 200]
[About input tables and joins](#) [page 201]
[About distinct rows on input tables](#) [page 209]
[About pre-filters and post-filters](#) [page 209]
[About mappings in the federation layer](#) [page 205]

11.3.1 Adding a federated table manually

Prerequisites

Before you begin, you must have a multisource-enabled data foundation based on at least one valid connection.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. From the **Federation Layer** pane, click **Add Federated Table**.
3. In the **Add Federated Table** dialog box, click the **Add Row** icon to add columns to the table.
4. For each column, edit the name, select a data type, and select whether or not input is needed.
Optionally, you can enter a description of the column.
5. Optionally, you can add a description for the federated table.
6. To save the federated table, click **OK**.
7. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

Next you need to define a mapping for the federated table just added.

Related Information

[How to build a data foundation](#) [page 138]
[About federated tables](#) [page 198]
[Adding a mapping](#) [page 207]

11.3.2 Adding a federated table from a data source

Prerequisites

Before you begin, you must have a multisource-enabled data foundation based on at least one valid connection.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. Click **Federation Layer** to open the *Federation Layer* pane, then click **Connections**.
3. From the **Connections** pane, select the table from the data source and drag it to the **Federation Data Flow** pane.
A federated table is automatically added with the same name and columns as the data source table. A default mapping is added that maps one-to-one the data source columns to the federated table columns.
4. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

Next you need to further define the input to the federated table by changing the mappings.

Related Information

[How to build a data foundation](#) [page 138]

[About federated tables](#) [page 198]

11.3.3 Editing a federated table

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, right-click the federated table that you want to edit and select **Edit**.
3. In the **Edit Federated Table** dialog box, you can perform the following operations:
 - Edit the table name.
 - Add or delete columns.

- Change the order of columns.
 - Edit column names and descriptions.
 - Change the data type of a column.
 - Change whether or not a column needs input.
 - Edit the table description.
4. To save the updates to the table, click **OK**.
 5. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About federated tables](#) [page 198]

11.4 About input tables and joins

Input tables define the input to a federated table. They can be data source tables or other federated tables.

The input tables are mapped to the federated table using mappings. In a mapping, a column in an input table is mapped to a column in the federated table. You can define the mapping formula so that one federated table column depends on one or more input table columns.

You can add multiple input tables to a mapping. In this case, you need to join the input tables.

To join input tables, first distinguish between core tables and non-core tables:

- Use a core table to choose the set of rows that will populate your federated table (the result set). When you set two or more tables as core, the result set is defined by the join of all the core tables. The core tables are joined using an inner join.
- Use non-core tables to extend the attributes of each row in the result set. A non-core table is joined using an outer join with core tables. If a row exists in the core table that doesn't match a non-core table row, it will return a row with nulls for the non-core columns.

The following restrictions apply to input tables and joins:

Direct joins between two non-core input tables are not allowed.

Cycles are not allowed (for example, if input table A is joined with B is joined with C, C cannot join to A.)

Note

When a table is a core table, the table name in the input tables pane in the mapping is in bold.

Related Information

[Adding input tables to a mapping](#) [page 202]

[Joining input tables](#) [page 203]

11.4.1 Adding input tables to a mapping

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table.
The mapping tabs appear in the Properties pane.
3. If the federated table has more than one mapping, select the appropriate mapping tab.
4. Do one of the following:

Option	Description
To add a data source table as an input table	Click Connections . Drag the data source table from the Connections pane to Input Tables in the mapping tab.
To add a federated table as an input table	Click Federation Layer . In the Properties tool bar, click the Add menu and select Add Input Table . You can also drag the federated table from the Federation Layer pane to Input Tables in the mapping tab.

5. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

You can now map the columns, edit mapping formulas, and add filters for the new mapping.

Related Information

[Mapping columns from the input table to columns of the federated table](#) [page 206]

[Editing a mapping formula](#) [page 206]

[About pre-filters and post-filters](#) [page 209]

11.4.2 Joining input tables

Context

Joining input tables applies when a mapping contains more than one input table.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table.
The mapping tabs appear in the **Properties** pane.
3. If the federated table has more than one mapping, select the appropriate mapping tab.
4. In the **Properties** pane, do one of the following:

Option	Description
Add a join visually	<p>This option lets you add a single-column join.</p> <p>In the input tables pane, click the column in the first input table and drag it to a column in the second input table.</p> <p>The two input tables are now joined by a default inner join on the selected columns.</p>
Add join with the join editor	<p>This option lets you add single-column and multi-column joins and simple formulas. For more information on the restrictions for join expressions, see the related topic.</p> <p>In the Properties tool bar, click the Add menu and select Add Join.</p> <p>In the Add Join dialog box, select a column in the left table and a column in the right table.</p> <p>You can edit the SQL for the join expression and click Validate to validate the SQL expression.</p> <p>To save the join definition, click OK.</p>

5. To select or unselect a table as a core table, right-click the table name in the input table pane and select **Core Table**.
When a table is a core table, the table name is bolded. For more information about core tables, see the related topic about input tables.
6. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

To edit a join, right-click the join line and select **Edit Join**.

Related Information

[About input tables and joins](#) [page 201]

11.4.3 Configuring meanings of input table joins using core tables

When you map multiple input tables to a federated table, you must distinguish between core tables and non-core tables.

- Use a core table to choose the set of rows that will populate your federated table (the result set).
When you set two or more input tables as core, the result set is defined by the join of all the core tables.
- Use non-core tables to extend the attributes of each row in the result set.

Example

The effect of setting an input table as core or non-core

Suppose that you have two input tables: **Customers** and **Orders**.

Setting on the Customers table	Setting on the Orders table	Result of a join between the two tables
core	non-core	all customers, including those who did not purchase anything (a left outer join)
core	core	only those customers who purchased something (an inner join)

The table below describes how you use core tables to configure meanings of input table joins:

Number and type of your input tables	Desired join result	Action
One input table	want to map some columns to the federated table	make sure the input table is a core table
Two input tables	want to display all values in all rows, including null values	make sure only one input table is a core table
Two input tables	want to display rows that contain null values	make sure both input tables are core tables
Three input tables	have a non-core table between two core tables	make sure you change the non-core table to a core table or one of the outer core tables to a non-core table

The effects on the federated table, of assigning an input table as a core table are represented in the following diagram (example in English):

Customer Table - Non-Core + Customer Address Table

Cust ID	Cust Name
1	N1
2	N2
3	N3
4	N4
5	N5

Cust ID	City	Zipcode
1	C1	Z1
2	C2	Z2
4	C4	Z4
5	C5	Null

-> Federated Table

Cust ID	Cust Name	City	Zipcode
1	N1	C1	Z1
2	N2	C2	Z2
3	N3	Null	Null
4	N4	C4	Z4
5	N5	C5	Null

Customer Table - Core + Customer Address Table

Cust ID	Cust Name
1	N1
2	N2
3	N3
4	N4
5	N5

Cust ID	City	Zipcode
1	C1	Z1
2	C2	Z2
4	C4	Z4
5	C5	Null

-> Federated Table

Cust ID	Cust Name	City	Zipcode
1	N1	C1	Z1
2	N2	C2	Z2
4	N4	C4	Z4
5	N5	C5	Null

11.5 About mappings in the federation layer

Mappings define transformations of values in your input tables and values in your federated tables.

When mapping columns, you need to know the data types of the columns that you are mapping. An icon showing the data type of the column is displayed in front of the column name. For example, **AB** indicates a string data type, and **12** indicates numeric. To see the data types of federated tables, you can also edit the table.

After you have mapped a column, you can edit the mapping formula to transform the value. For example, you can use formulas to construct new values in the federated table column, combine multiple values, or calculate results.

Related Information

[Mapping columns from the input table to columns of the federated table](#) [page 206]

[Editing a mapping formula](#) [page 206]

[About pre-filters and post-filters](#) [page 209]

[About input tables and joins](#) [page 201]

[Adding a mapping](#) [page 207]

[Activating and deactivating mappings](#) [page 208]

11.5.1 Mapping columns from the input table to columns of the federated table

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table.
The mapping tabs appear in the **Properties** pane.
3. If the federated table has more than one mapping, select the appropriate mapping tab.
4. Select a column in an input table and drag it to a column in the federated table.
A mapping line appears between the columns.

Next Steps

Edit the mapping formula for the column.

Related Information

[Editing a mapping formula](#) [page 206]

[About mappings in the federation layer](#) [page 205]

11.5.2 Editing a mapping formula

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table.
The mapping tabs appear in the **Properties** pane.
3. If the federated table has more than one mapping, select the appropriate mapping tab.
4. Right-click the column in the federated table and select **Edit Mapping Formula**.
5. Edit and validate the SQL expression for the mapping formula in the SQL Expression Editor, and when you are finished, click **OK**.
6. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About the SQL Expression Editor](#) [page 207]




11.5.3 About the SQL Expression Editor

The SQL Expression Editor helps you write valid SQL expressions.

You can type SQL directly into the [Expression](#) box, drag and drop table or column names from the **Sources** pane, and drag and drop operators and database functions from the **Functions** pane. These panes are described in the following table. To display a resource pane, click the icon in the tool bar of the **Expression** pane.

Click the **Validate** icon in the tool bar of the **Expression** pane to check if the expression you have defined is valid SQL.

To save the expression, click **OK**.

Icon	Description
 Sources	The list of tables and columns in the data foundation. To see a list of values for a column, click the  icon next to column name.
 Functions	The list of functions that can be used in the expression. The functions are grouped by type: <ul style="list-style-type: none">• Operators: Common database operators, for example *, SUM, IS NOT NULL.• Database Functions: The SQL functions that are valid for multisource-enabled data foundations. See the related topic about SAP BusinessObjects SQL functions.

Related Information

[SAP BusinessObjects SQL function reference for multisource-enabled universes](#) [page 360]

11.5.4 Adding a mapping

Prerequisites

If you have not yet defined a default mapping for the federated table, see the related topic about mappings. This task describes adding mappings in addition to the default mapping.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table for which you want to add a mapping.
3. In the **Properties** pane, click the **Add Mapping** tab.
4. Enter a name for the mapping and click **OK**.
5. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

You can now add input tables, map the columns, edit mapping formulas, and add filters for the new mapping.

Related Information

[About mappings in the federation layer](#) [page 205]

[Adding input tables to a mapping](#) [page 202]

[Mapping columns from the input table to columns of the federated table](#) [page 206]

[Editing a mapping formula](#) [page 206]

[About pre-filters and post-filters](#) [page 209]

11.5.5 Activating and deactivating mappings

Context

The effective mapping for a federated table is the implicit union of all activated mappings.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table in the **Federation Data Flow** pane.
3. Right-click the mapping that you want to activate or deactivate and select **Activate**.
When a mapping is deactivated, the mapping name is crossed-out in the table view.
4. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

11.6 About distinct rows on input tables

The distinct rows feature lets you specify whether the rows coming from an input table should be distinct. You can turn on the distinct rows feature on each input table.

11.6.1 Activating and deactivating distinct rows

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table.
The mapping tabs appear in the Properties pane.
3. If the federated table has more than one mapping, select the appropriate mapping tab.
4. Right-click the input table name and select **Distinct Row**.

Note

If there is a checkmark before the Distinct Rows menu item, it indicates that the feature is active and if there is not one, it means the feature is not active.

5. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

11.7 About pre-filters and post-filters

Filters let you transform data in mappings in two ways:

- Pre-filters let you limit the source data that is queried in the mapping. For example, you can use a filter to limit customer data to those customers who are born after a certain date.
You can use a pre-filter on each input table that is used in a mapping.
- Post-filters let you limit the data after it has been treated by table joins. Use post-filters when the filter definition depends on columns from more than one input table. For example, to limit the orders to customers who were 18 years old or older on the order date.
You can use one post-filter per mapping.

Pre-filters are applied before the table joins. Post-filters are applied after the table joins. Mapping formulas are applied after the post-filters.

11.7.1 Adding and editing pre-filters

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table.
The mapping tabs appear in the **Properties** pane.
3. If the federated table has more than one mapping, select the appropriate mapping tab.
4. Select the input table and do one of the following:

Option	Description
To add a pre-filter	Right-click the input table name and select Add Pre-filter .
To edit and existing pre-filter	Right-click the input table name and select Edit Pre-filter .

Note

One pre-filter is allowed per input table.

5. Edit and validate the SQL expression for the pre-filter in the SQL Expression Editor, and when you are finished, click **OK**.
6. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About the SQL Expression Editor](#) [page 207]

[About pre-filters and post-filters](#) [page 209]

11.7.2 Editing post-filters

Prerequisites

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, select the federated table for which you want to add a mapping.

3. If the federated table has more than one mapping, select the appropriate mapping tab.

4. Click the **Edit Post-Filter** icon .

Note

One post-filter is allowed per mapping.

5. Edit and validate the SQL expression for the post-filter in the SQL Expression Editor, and when you are finished, click **OK**.

6. To save the data flow in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[About the SQL Expression Editor](#) [page 207]

[About pre-filters and post-filters](#) [page 209]

11.8 Showing values in a federated table

Context

The show values command applies the pre-filters, joins, post-filters, and mapping formulas. If the federated table contains an input column, you are prompted for a value.

By default, show values opens a tab in the editor to display the values. You can set a preference to have the values open in a dedicated view or a dialog box. For more information, see the related topic.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. In the **Federation Layer** pane, right-click the federated table and select **Show Table Values**.

The show values window appears. To see what you can do in this window, see the related topic on showing values in a data source.

Related Information

[Showing values in a data source](#) [page 185]

[Setting preferences for showing values](#) [page 34]

11.9 Checking integrity of the federation layer

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. Do one of the following:

Option	Description
To check integrity of the federation layer only	Click Federation Layer
To check integrity of the data foundation, including the federation layer	Click Data Foundation and select the top level of the data foundation in the tree view.

3. From the main tool bar, select the **Check Integrity** icon .

Results

For more information about check integrity rules and the results of an integrity check, see the related topic.

Related Information


[Running check integrity](#) [page 316]

11.10 Inserting a federated table into the data foundation

Prerequisites

Before you can insert a federated table into the data foundation, you must define the federated table and data flow in the federation layer.

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. Select **Insert Federated Table** from the **Insert**  menu in the Data Foundation view.
3. In **Insert Federated Tables**, select the federated tables you want to insert, and click **OK**.

4. Save the data foundation by clicking the **Save** icon in the main tool bar.



Related Information

[Building the federation data flow](#) [page 196]

11.11 Refreshing the structure of the federation layer

Procedure

1. Open the data foundation by double-clicking the data foundation name in the Local Projects View.
2. Do one of the following:

Option	Description
To refresh the structure of the federation layer only	Click Federation Layer . In the <i>Federation Layer</i> pane, click the Refresh Structure icon  .
To refresh the structure of the data foundation, including the federation layer	Click Data Foundation and then select Actions > Refresh Structure  .

For the Federation Layer, the wizard detects the following changes and lists them each in their own dialog box. In each case, you select which of the proposed changes to make in the federation layer.

- Tables in the federation layer that were deleted in the database. The wizard proposes to delete these tables and any related joins from the federation data flow.
 - Columns in federation layer tables that were deleted in the database tables. The wizard proposes to update each corresponding table in the federation layer to delete these columns and the joins that use these columns.
 - Columns added in the database. The wizard proposes to update each corresponding table in the federation layer to add these columns.
 - Column data types changed in the database. The wizard proposes to update the data type of each column in the data foundation that is different from the database column type.
3. Select changes in a summary dialog box and click **Finish** to continue with the refresh.
 4. To save the changes in the federation layer, save the data foundation by clicking the **Save** icon in the main tool bar.

Next Steps

Check the integrity of the federation layer to detect definitions in your data flow that need to be updated because of the changes made by the refresh structure.

Related Information

[Checking integrity of the federation layer](#) [page 212]

12 Working with business layers

12.1 About business layers

A business layer is a collection of metadata objects that map to SQL or MDX definitions in a database, for example, columns, views, database functions, or pre-aggregated calculations. The metadata objects include dimensions, hierarchies, measures, attributes, and predefined conditions. Each object corresponds to a unit of business information that can be manipulated in a query to return data. Business layers can be created directly on an OLAP cube, or on a data foundation that is built on a relational database.

When a business layer is complete, it is published to a repository or a local folder as a universe. A universe is a published .unx file that includes a business layer and its connection to an OLAP cube, or a business layer and its corresponding data foundation. The universe is available in the repository to SAP BusinessObjects data analysis and report creation applications.

The principle role of the business layer is to define and organize metadata before it is published as a universe. An alternative way to understand the business layer is to think of it as a metadata workbench that a designer uses to assemble and modify a metadata set before publishing as a universe for data analysis and report creation applications.

Related Information

[About business layer objects](#) [page 215]

[About business layer properties](#) [page 227]

[How to build a relational business layer](#) [page 218]

[How to build an OLAP business layer](#) [page 222]

12.2 About business layer objects

The **Business Layer** objects pane contains the metadata objects that make up the business layer.








Every object in the business layer has a name that can be edited. When naming objects, use business vocabulary familiar to the users who will be using the universe for querying, analysis, and reporting.







Objects can have three states:

- **Active:** Object is visible in the Query Panel. This is the default state.
- **Hidden:** Object is valid but not available in the Query Panel (used by other objects as a hidden object).
- **Deprecated:** Object is hidden and not valid. One possible use for this state is when the target database field no longer exists, but you want to keep the object for possible future use.

Each object in the business layer has properties that are applied in the published universe. You can define properties when you insert an object, and modify object properties at any time. See the related topics about inserting and editing objects.

Depending on the type of data source, you can create and edit the following types of objects in the business layer:

Object	Description
 Folder	A folder is a container that holds a group of related objects. You create folders to group objects that have a common purpose in the business layer. The folder has no role in a query. It is only used to organize objects.
 Dimension	<p>A dimension is an object that maps to one or more table columns or a function in a database and represents an axis of analysis in a query. For example, Product, Geography, Time, and Employee are common dimensions. Each dimension classifies an aspect of an activity in a business environment.</p> <p>In a business layer, dimensions represent contextual information (the axes of analysis).</p>
 Measure	<p>Measures are objects that represent calculations and aggregate functions that map to statistical and analytic data in the database.</p> <p>In a business layer, measures represent the factual information (data).</p> <p>Numeric data is usually, but not always, the source of a measure. Aggregating the information must make sense for the object to be a measure. For example, summing up sales revenue makes sense, so Sales Revenue is a measure. Summing product list prices is not necessarily useful, so List Price is a dimension, or perhaps an attribute of the Product Dimension.</p> <p>You can create measures from non-numeric objects by counting things. This can result in measures like Number of Orders.</p>
 Attribute  Measure attribute	<p>An attribute is an object attached to a parent object that provides additional descriptive information about the parent. Attributes can be defined for dimensions, measures, hierarchies, and levels.</p> <p>In an OLAP business layer, a measure attribute provides information for the formatted value.</p>
 Filter	<p>A filter is a condition object that limits the data returned in a query. Filters can be inserted into the Query Filters pane in the Query Panel to be applied to the query.</p> <p>Native filters are defined by an SQL WHERE clause on data foundation tables. Native filters apply to business layers based on data foundations.</p> <p>Business filters are defined by creating and combining conditions on dimensions and measures in the business layer.</p>
 Analysis dimension (OLAP only)	An analysis dimension allows you to logically group dimensions and hierarchies that share the same axis of analysis. Analysis dimensions are often used for hierarchical analysis.

Object	Description
	<p>You define a default hierarchy for the analysis dimension. This is the hierarchy used when the entire analysis dimension is included as a result object in a query. Default hierarchies have the following icon: .</p>
 Hierarchy (OLAP only)	<p>A hierarchy is the representation in the business layer of the hierarchy in the OLAP cube. If the hierarchy in the cube is level-based, level objects in the business layer represent the levels.</p> <p>If the hierarchy in the cube is value-based (parent-child), the levels are not represented in the business layer. The levels are visible when previewing members and in the Member Selector. Value-based hierarchies generated automatically in the business layer have the following icon: .</p>
 Level (OLAP only)	<p>A hierarchy level in a level-based hierarchy.</p>
 Named set (OLAP only)	<p>A named set is a collection of members of a hierarchy in the business layer.</p> <p>A native named set is defined using an MDX expression. For some connections, native named sets are created automatically to represent named sets in the cube.</p> <p>A business named set is defined by selecting members.</p>
 Calculated member (OLAP only)	<p>A calculated member is a member of a hierarchy that is calculated using an explicitly defined MDX expression that can include data from the OLAP cube, mathematical operators, numbers, and functions.</p> <p>Calculated members are available in the Member Selector when creating queries.</p>

Related Information

[Working with business layer objects](#) [page 243]

[Inserting a folder](#) [page 243]

[Inserting and editing dimensions](#) [page 244]

[Inserting and editing measures](#) [page 249]

[Inserting and editing attributes](#) [page 253]

[Inserting and editing filters](#) [page 255]

[Inserting and editing analysis dimensions](#) [page 257]

[Inserting and editing hierarchies](#) [page 258]

[Inserting and editing hierarchy levels](#) [page 260]

[Inserting and editing named sets](#) [page 261]

[Inserting and editing calculated members](#) [page 262]

12.3 How to build a relational business layer

Prerequisites

Before you begin:

- You need a project in the Local Projects View.
- You need a data foundation saved in the same local project.

Context

For links to more detailed information on each step, see the Related Topics.

Procedure

1. To start the *New Business Layer* wizard, do one of the following:

- Right-click a data foundation in the Local Projects View and select ► **New** ► **Business Layer** ►.
- Right-click the project folder in the Local Projects View and select ► **New** ► **Business Layer** ► and select **Relational data foundation** as the data source.

The business layer is created in a .blx file in the local project folder. It opens automatically in the Business Layer Editor.

2. Build the business layer:

If you created the business layer with the **Automatically create folders and objects** option selected (default), all objects are created in the business layer as dimensions. You need to specify the measures explicitly using the **Turn into Measure with Aggregation Function** command.

Otherwise, insert the business layer objects:

- a) Insert folders and subfolders to organize the business layer.

Note

When you drag a table into the business layer from the data foundation, a folder is automatically inserted.

- b) Drag and drop tables and columns into the desired folders and rename the objects if needed.
 - c) Specify the measures using the **Turn into Measure with Aggregation Function** command.
3. You can enhance the function of the business layer in several ways, for example:
 - Insert attributes to provide descriptive information for dimensions and measures
 - Insert additional measures
 - Insert pre-defined filters (mandatory or optional) that can limit data returned in queries

- Insert parameters with optional prompts
 - Insert custom lists of values to be associated with a prompt
 - Insert navigation paths to define drill paths
 - Create business layer views to restrict the objects seen in the Query Panel
 - Set SQL options and SQL generation parameters in the business layer properties
 - Set up aggregate awareness to improve query performance
4. Run an integrity check to validate the dependencies, object expressions, and parameters and lists of values. Right-click the business layer name in the **Business Layer** pane, and select **Check Integrity**.
 5. Save the business layer by clicking the **Save** icon in the main tool bar.

Next Steps

Some commands to help you maintain the business layer are listed below.

- If you change object definitions, use **Show Local Dependencies** to find the other business layer objects and data foundation objects that might be impacted by the changes.
- If the related data foundation is modified, you must refresh the business layer manually. The related topic describes commands to help you do this.
- Use **Change Data Foundation** to change the source data foundation for the business layer.
- For multisource-enabled business layers, use **Compute Statistics** to improve query performance.

Related Information

[Creating a local project](#) [page 79]

[How to build a data foundation](#) [page 138]

[About the Business Layer Editor](#) [page 225]

[About business layer objects](#) [page 215]

[Turning a dimension or attribute into a measure](#) [page 251]

[Inserting a folder](#) [page 243]

[Inserting and editing dimensions](#) [page 244]

[Inserting dimensions directly from the data foundation](#) [page 246]

[Inserting and editing measures](#) [page 249]

[Inserting and editing filters](#) [page 255]

[Inserting and editing a parameter](#) [page 281]

[Inserting or editing a list of values](#) [page 284]

[Associating a list of values with a business object](#) [page 287]

[Associating a list of values with a prompt defined in the business layer](#) [page 287]

[About business layer views](#) [page 279]

[About business layer properties](#) [page 227]

[About aggregate awareness](#) [page 241]

[About resource dependencies](#) [page 318]

[Running check integrity](#) [page 316]

[About refreshing business layers](#) [page 291]

[Changing the data source of a business layer](#) [page 231]

[About computing statistics for optimized query execution](#) [page 294]

12.3.1 Specifying the type of data source for a business layer

Context

This section describes the [Select the type of data source for the business layer](#) page of the New Business Layer wizard.

You choose to create a business layer from either a relational or an OLAP data source.

Data Source Type	Description
Relational	The business layer is based on a data foundation. You can select any data foundation in the current project folder.
OLAP	The business layer is based on an OLAP cube. You can select any OLAP connection or connection shortcut in the current project folder.

Procedure

1. Click one of the data source types in the list.
2. Click **Next**.

12.3.2 Naming a business layer

Context

This section describes the [Resource Name](#) page of the New Business Layer wizard.

You enter a name and description of the business layer. This is the name of the universe that is published from the business layer.

Related Information

[Selecting a data foundation for a business layer](#) [page 221]

[Selecting an OLAP connection and cube for a business layer](#) [page 224]

12.3.3 Selecting a data foundation for a business layer

Context

This section describes the *Select Data Foundation* page of the New Business Layer wizard.

Select a data foundation as the data source for the new business layer. You can choose to do one of the following:

- Automatically create the business layer objects from the tables and columns in the data foundation.
- Create an empty business layer. You must manually add the objects from the data foundation after creation.

Procedure

1. Click the browse button at the end of the Data Foundation text field.
A list of available data foundations appears.
2. Select a data foundation in the list and click **OK**.

The data foundation name appears in the name field. The **Automatically create folders and objects** option is selected by default.

Business layer object names are generated based on table and column names (except for data foundations based on SAP ERP and SAP NetWeaver BW, which use a dedicated strategy for naming objects). You can set an application preference to determine how names are generated. For more information, see the related topic about setting preferences for the Business Layer Editor.

3. Do one of the following:
 - If you want to automatically populate the business layer with objects and folders, click **Finish**.
 - If you do not want to automatically populate the business layer, unselect the option and click **Finish**.

Results

The new business layer opens in an editing tab. You can now insert and edit the business layer objects.

Next Steps

The **Automatically create folders and objects** option creates all objects in the business layer as dimensions. Specify the measures explicitly using the **Turn into Measure with Aggregation Function** command. For more information, see the related topic.

Related Information

[How to build a relational business layer](#) [page 218]

[Turning a dimension or attribute into a measure](#) [page 251]

[About business layer objects](#) [page 215]

[Setting preferences for the Business Layer Editor](#) [page 25]

12.4 How to build an OLAP business layer

Prerequisites







Before you begin:

- You need a project in the Local Projects View.
- You need an OLAP connection or connection shortcut saved in the same local project.

Context

For links to more detailed information on each step, see the Related Topics.

Procedure

1. To start the [New Business Layer](#) wizard, do one of the following:
 - Right-click an OLAP connection or connection shortcut in the Local Projects View and select  **New**  **Business Layer** .
 - Right-click the project folder in the Local Projects View and select  **New**  **Business Layer**  and select **OLAP connection** as the data source.

The business layer is created in a .blx file in the local project folder. It opens automatically in the Business Layer Editor.

2. Objects in the business layer are inserted automatically based on the cube. You can enhance the function of the business layer in several ways, for example:
 - Insert analytical dimensions, hierarchies, and attributes
 - Insert named sets
 - Insert calculated members
 - Insert measures and their formatted value attributes
 - Insert pre-defined filters (mandatory or optional) to limit data returned in queries
 - Insert parameters with optional prompts

- Insert custom lists of values to be associated with a prompt
- Create business layer views to restrict the objects seen in the Query Panel
- 3. Run an integrity check to validate the dependencies, object expressions, and parameters and lists of values. Right-click the business layer name in the **Business Layer** pane, and select **Check Integrity**.
- 4. Save the business layer by clicking the **Save** icon in the main tool bar.

Next Steps

Some commands to help you maintain the business layer are listed below.

- If you change object definitions, use **Show Local Dependencies** to find the other business layer objects that might be impacted by the changes.
- If the underlying data source is modified, use **Refresh Structure** to refresh the business layer.
- Use **Change OLAP Connection** to change the connection for the business layer and to edit the OLAP data source properties.

Related Information

[Creating a local project](#) [page 79]

[Creating an OLAP connection](#) [page 123]

[About business layer objects](#) [page 215]

[Inserting and editing analysis dimensions](#) [page 257]

[Inserting and editing hierarchies](#) [page 258]

[Inserting and editing hierarchy levels](#) [page 260]

[Inserting and editing attributes](#) [page 253]

[Inserting and editing named sets](#) [page 261]

[Inserting and editing calculated members](#) [page 262]

[Inserting and editing measures](#) [page 249]

[Inserting and editing filters](#) [page 255]

[Inserting and editing a parameter](#) [page 281]

[Inserting or editing a list of values](#) [page 284]

[Associating a list of values with a business object](#) [page 287]

[Associating a list of values with a prompt defined in the business layer](#) [page 287]

[About business layer views](#) [page 279]

[Running check integrity](#) [page 316]

[Refreshing an OLAP business layer](#) [page 293]

[Changing the data source of a business layer](#) [page 231]

12.4.1 Selecting an OLAP connection and cube for a business layer

Context

This section describes the [Select OLAP Connection](#) page of the New Business Layer wizard.

Select an OLAP connection and the OLAP cube as the data source for the new business layer.

Note

You cannot build a business layer on an **SAP BICS Client** connection although these connections appear in the list. Use **SAP BICS Client** connections in SAP BusinessObjects query and reporting applications to connect directly to the BEx query. No business layer or universe is required to access BEx queries. For information on building a universe on an SAP NetWeaver BW connection, see the related topic.

OLAP connection options	Description
OLAP connection	Click the browse button at the end of the text field to select an OLAP connection or connection shortcut defined in the project.
Detect measure projection function	If not selected the database delegated function is applied.
Create attribute from unique name	An attribute is created for the unique name for each dimension.
Search	Enter a search string for a cube and click the search icon.
List of connection cubes	List of available cubes to the connection. If there are multiple cubes, you browse to and select the target cube.

Related Information

[Selecting objects from an OLAP cube for a business layer](#) [page 225]

[About projection functions](#) [page 251]

[Using SAP NetWeaver BW data sources](#) [page 40]

12.4.2 Selecting an Essbase Accounts dimension

This section describes the [Select Accounts Dimension](#) page of the New Business Layer wizard.

For connections to Essbase data sources, the New Business Layer wizard creates measures in the business layer from the objects in the specified Accounts dimension in the data source.

Select one dimension in the list you want to use as the Accounts dimension and click **Next**.

12.4.3 Selecting objects from an OLAP cube for a business layer

Context

This section describes the [Select Objects](#) page of the New Business Layer wizard.

Expand the object nodes under the selected cube and select the objects to be included in the new business layer. Click **Finish** when you have completed the selection.

The new business layer appears in the Business Layer pane.

12.5 About the Business Layer Editor

You use the Business Layer Editor to create and edit business layer objects and properties. This topic describes how to navigate the editor. For steps to help you build the structure of your business layer, see [How to build a relational business layer](#) [page 218] or [How to build an OLAP business layer](#) [page 222].

The Business Layer Editor is divided into browsing panes on the left, an editing pane on the upper right, and a data source pane on the lower right.

The browsing panes allow you to work with different elements of the business layer. Access the panes by clicking the corresponding tab:

- **Business Layer**
- **Queries**
- **Parameters and Lists of Values**
- **Navigation Paths**

For more information about what you can do in each of the browsing panes, see the related topic.

The **Business Layer** is the default browsing pane. It displays a tree view of the objects in the business layer. The following options are available for displaying and navigating the business layer tree view:

- Filter by business layer view
- Search for an object
- Change display options: Show or hide objects, show unique names

The editing pane lets you edit the properties of the object or element selected in the browsing pane.

The data source pane displays the data foundation or the OLAP connection information:

- The data foundation master view containing all tables and joins displays by default. Tabs for other data foundation views, if defined, appear at the bottom of the data source pane. To change views, click the tab.
- The OLAP metadata in the connection displays in the left side of the data source pane. Select a metadata object to display its properties in the right side of the pane.

Related Information

[About business layer properties](#) [page 227]

[About business layer objects](#) [page 215]

[About queries in a business layer](#) [page 290]

[About parameters](#) [page 280]

[About lists of values](#) [page 283]

[About navigation paths for objects](#) [page 288]

[About business layer views](#) [page 279]

[Filtering by business layer view](#) [page 280]

[Searching for business layer objects](#) [page 278]


[Changing display options of the business layer tree view](#) [page 226]

12.5.1 Changing display options of the business layer tree view

Context

When editing a business layer, the **Business Layer** browsing pane displays a tree view of the objects in the business layer. Use this procedure to change the display mode of the business layer objects.

Procedure

1. Click the **Display options** icon  at the top of the **Business Layer** browsing pane.
2. For business layers based on an OLAP connection, select one of three options:
 - **Display caption** to display the object names.
 - **Display unique name** to display the unique object name from the cube.
 - **Display caption and unique name**
3. To show only active objects in the business layer tree view, select **Hide non-active objects**.

Results

The display options remain in effect until you close the editor.

Related Information

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

12.6 About business layer properties

The following properties and options are defined for the entire business layer. The restrictions are applied in the published universe.

Property		Description
Name		Identifies the business layer and also the universe when the business layer is published.
Description		Describes the universe purpose and content. This description is available to display in the query and reporting tools that use the published universe.
Query Limits	Limit result set size to	Specifies the number of rows that are returned in a query. This limits the number of rows returned, but does not restrict the RDBMS from processing all rows in the query. It only limits the number once the RDBMS has started to send rows.
	Limit execution time to	Specifies the number of minutes to limit the time passed for query execution, but does not stop the process on the database.
	Warn if estimate exceeds	When selected, displays a message when an estimation of the execution time exceeds the specified number of minutes.
Query Options (apply to business layers based on data foundations)	Allow use of subqueries	When selected, allows subqueries in a query.
	Allow use of union, intersect, and minus operators	When selected, lets you combine queries using the data set operators union, intersect, and minus, to obtain one set of results.
	Allow complex operands in Query Panel	When selected, displays complex operands in the list of operands available when defining a filter in the Query Panel.
	Multiple SQL statements for each measure	When selected, one SQL query is generated for each measure or group of measures belonging to a different fact table, or for measures that have a WHERE clause (filtered measure). If the measure objects are based on columns in the same table, then the separate SQL queries are not generated, even if this option is selected.
	Allow query stripping	When selected, a report user can allow query stripping to be enabled for relational universes. (Query stripping is enabled by default in OLAP universes).

Property		Description
		<p>Query stripping is used only by SAP BusinessObjects Web Intelligence.</p> <p>For a description of query stripping and how it can impact reporting in relational universes, see the related topic.</p>
Data source		<p>Specifies the data source of the business layer: either a data foundation or an OLAP connection.</p> <p>The Change Data Foundation button lets you change the underlying data foundation.</p> <p>The Change OLAP Connection button lets you change to OLAP connection and edit data source properties.</p>
SQL Parameters (apply to business layers based on data foundations)		Specifies custom values for SQL generation parameters that override the default values or any customized value in the data foundation properties.
Comments		Contains comments about the business layer.
Summary		Displays a summary of the number of each type of object defined in the business layer. For business layers based on a data foundation, the type and number of data foundation objects are also displayed.

Related Information

[Editing the business layer name, description, and comments](#) [page 230]

[Editing query limits and options in the business layer](#) [page 231]

[Changing the data source of a business layer](#) [page 231]

[About query stripping](#) [page 229]



[Setting SQL generation parameters in the business layer](#) [page 232]

[Displaying a business layer summary](#) [page 233]

12.6.1 OLAP data source properties

The following properties apply to the OLAP data source for the business layer:

Property	Description
OLAP Connection	The connection or connection shortcut that provides the access to the OLAP data source.

Property	Description
	To change the connection, click the browse icon  at the end of the field to open a list of available connections.
Cube	<p>The cube selected for the current connection. You can select a different cube only if a cube was not specified when the connection was defined.</p> <p>To change the cube, click the browse icon  at the end of the field to open a list of available cubes.</p>
Accounts Dimension	<p>For connections to Essbase data sources, the dimension in the data source to use as the Accounts dimension. Select a dimension from the list.</p> <p>When refreshing the structure of the business layer, measures are created in the business layer from the objects in the specified Accounts dimension.</p>
END_MDX value	<p>The value of the END_MDX parameter.</p> <p>The END_MDX parameter is equivalent to the END_SQL parameter available for universes based on data foundations. The value of the END_MDX is added to the end of every MDX statement.</p> <p>For example, you can use the END_MDX parameter to track the database server activity by tracing who is running queries. The solution consists of adding a comment at the end of every MDX query with information about the user and the universe. For example:</p> <pre>//User: @Variable('BOUSER') Universe: @Variable('UNVNAME')</pre>

12.6.2 About query stripping

Query stripping is a reporting feature that can be used to optimize performance. Query stripping is used only by SAP BusinessObjects Web Intelligence.

For relational universes, query stripping is only enabled if the following parameters are set:

- The **Allow query stripping** option is selected in the business layer properties in the information design tool (unselected by default).
- The **Enable query stripping** option is selected for the data provider in the Query Properties in Web Intelligence.
- The **Enable query stripping** option is selected in the Document Properties in Web Intelligence (selected by default if query stripping is enabled for the data provider).

For OLAP universes, query stripping is enabled by default.

When query stripping is enabled, the query is rewritten to reference only objects that are used in the report. Let's take for example a query that contains three result objects: **Country**, **City**, and **Revenue**. A report based on this query may contain only **City** and **Revenue**. If query stripping is enabled, when the report is refreshed, in most cases the query will only retrieve the data for **City** and **Revenue**.

In relational universes, a report with query stripping enabled may return different data than when query stripping is disabled, depending on the schema of the data foundation. Let's look again at the example of a query that

contains **Country**, **City**, and **Sales Revenue**. In the data foundation, there is a self-restricting join on the **Country** table that restricts the country to the US. With query stripping disabled, the report on **City** and **Revenue** returns revenue for only cities in the US. With query stripping enabled, the report returns revenue for cities in all countries, because the **Country** table was stripped out of the query.

Enhanced Query Stripping

The `USE_ENHANCED_QUERY_STRIPPING` parameter allows you to take advantage of enhancements to the query stripping method. Normal query stripping rewrites the query to contain only the objects referenced in the report and any joins concerned by those objects. Enhanced query stripping only optimizes the `SELECT` and `GROUP BY` clauses to avoid fetching unused data, but it doesn't modify the other clauses or the joins. It is recommended to use enhanced query stripping in the following situations:

- The data foundation contains outer joins.
- The data foundation contains self-restricting joins (column filters).
- The data foundation contains shortcut joins.

If aggregate awareness is defined in the business layer (using the `@Aggregate_aware` function in the definition of business layer objects), enhanced query stripping is used in every case, even if the `USE_ENHANCED_QUERY_STRIPPING` parameter is not set.

The `USE_ENHANCED_QUERY_STRIPPING` parameter is not set by default. It can be set in the data foundation or business layer. For more information, see the related topics.

Related Information

[Editing query limits and options in the business layer](#) [page 231]

[USE_ENHANCED_QUERY_STRIPPING](#) [page 445]

[Setting SQL generation parameters in the data foundation](#) [page 182]

[Setting SQL generation parameters in the business layer](#) [page 232]

12.6.3 Editing the business layer name, description, and comments

Procedure

1. Open the business layer in the editor by double-clicking the business layer in the Local Projects View.
2. Make sure the top level of the business layer is selected in the tree view in the **Business Layer** pane.
3. Change the business layer properties in the editing pane:
 - To change the business layer name, edit **Name**.
 - To enter or edit the business layer description, click the **Properties** tab.

- To enter or edit comments for the business layer, click the **Comments** tab.
4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer properties](#) [page 227]

12.6.4 Editing query limits and options in the business layer

Procedure

1. Open the business layer in the editor by double-clicking the business layer in the Local Projects View.
2. Make sure the top level of the business layer is selected in the tree view in the **Business Layer** pane.
3. Click the **Query Options** tab in the editing pane.
4. Select or unselect options and edit limit values as required. For a description of the options, see the related topic.
5. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer properties](#) [page 227]

12.6.5 Changing the data source of a business layer


Prerequisites

To change the data source for a business layer, the new data source (data foundation, OLAP connection or connection shortcut) must be saved in the same local project folder as the business layer.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Make sure the top level of the business layer is selected in the tree view in the **Business Layer** pane.
3. Make sure the **Properties** tab is selected in the editing pane.

4. Do one of the following depending on the type of data source for the business layer:

Option	Command
For data foundation sources	Click Change Data Foundation . Select the new data foundation from the list and click OK .
For OLAP sources	<p>Click Change OLAP Connection.</p> <p>In the <i>Edit OLAP Data Source Properties</i> dialog box, click the browse icon  at the end of the OLAP Connection text box. Select the new OLAP connection or connection shortcut and click OK.</p> <div><p>i Note</p><p>For more information on advanced OLAP properties, see the related topic.</p></div>

5. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[OLAP data source properties](#) [page 228]

12.6.6 Setting SQL generation parameters in the business layer

Context

The custom values for SQL generation parameters in the business layer override the default values or any customized value set in the data foundation properties.

Procedure

1. Open the business layer in the editor by double-clicking the business layer in the Local Projects View.
2. Make sure the top level of the business layer is selected in the tree view in the **Business Layer** pane.
3. Make sure the **Properties** tab is selected in the editing pane.
4. Click the **Parameters** button.
5. In the *Query Script Parameters* dialog box, edit the parameters:
The currently defined SQL generation parameters are listed. Non-default parameters and parameters with non-default values are in bold type.

Option	Action
Change the value of an existing parameter.	Click the Value column and select or enter the new value.
Add a predefined parameter.	To display the list of predefined parameters, click the arrow in the list box next to the Add button. Select the parameter from the list and click Add .
Add a custom parameter.	Make sure no predefined parameter is listed in the box next to the Add button, then click Add . A parameter with a default name is added to the table. To edit the parameter name, click the Name column. Click the Value column to enter a value.

To see a description of all the predefined SQL generation parameters and their values, click the help button.

6. To return to the default list of parameters and the default values, click **Default Values**. This removes any added parameters from the list and sets all values to the default.
7. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About SQL Generation Parameters](#) [page 433]

12.6.7 Displaying a business layer summary

Context

Use this command to display the number of each type of object defined in the business layer. For business layers based on a data foundation, the type and number of data foundation objects are also displayed.

Procedure

1. Open the business layer in the editor by double-clicking the business layer in the Local Projects View.
2. Make sure the top level of the business layer is selected in the tree view in the **Business Layer** pane.
3. Make sure the **Properties** tab is selected in the editing pane.
4. Click the **Summary** button.

Results

The business layer summary displays in a new dialog box.

12.7 About index awareness

In a relational business layer, index awareness is the ability to take advantage of the indexes on key columns to improve query performance.

The objects in the business layer are based on database columns that are meaningful for querying data. For example, a Customer object retrieves the value in the customer name column of the customer table. In many databases, the customer table has a primary key (for example an integer) to uniquely identify each customer. The key value is not meaningful for reporting, but it is important for database performance.

When you set up index awareness, you define which database columns are primary and foreign keys for the dimensions and attributes in the business layer. The benefits of defining index awareness include the following:

- Joining and filtering on key columns are faster than on non-key columns.
- Fewer joins are needed in a query, therefore fewer tables are requested. For example, in a star schema database, if you build a query that involves filtering on a value in a dimension table, the query can apply the filter directly on the fact table by using the dimension table foreign key.
- Uniqueness in filters and lists of values is taken into account. For example, if two customers have the same name, the application retrieves only one customer unless it is aware that each customer has a separate primary key.

For details on how to define primary and foreign keys on business layer objects, see the related topic.

Related Information

[Defining keys for dimensions and dimension attributes](#) [page 247]

[Creating an index-aware prompt](#) [page 282]

12.8 About analytic functions

You use the information design tool to define analytic functions for objects in a universe in order to calculate, for example, rankings, moving sums or averages, and relative calculations.

An analytic function is a function provided by the relational database that performs an analytical task on a result set. An analytic function in a query returns, with each row in the result set, a calculation from a group of rows. The groups of rows can be ordered and partitioned.

For example, you can use analytic functions to retrieve the following results:

- The rank of a record, for example the rank of retail stores by the sales amount for last month.
- Moving sum or average, for example the average volume of sales in a three-month period.
- Displaying the same information in different contexts, for example this quarter's sales and last quarter's sales.
- Relative calculations, for example the difference between the sales of this quarter and the highest sales amount ever.

Following are some examples of analytic functions:

- Aggregation-like functions: SUM, COUNT, AVG, STDDEV, MEDIAN, VARIANCE
- Order-based functions: RANK, PERCENT_RANK, DENSE_RANK, LEAD, LAG, FIRST_VALUE, ROW_NUMBER

For a full description of the analytic functions available in your database, refer to the database documentation.

In the information design tool you can use analytic functions in the SELECT statement for measures and dimensions in the business layer, and for derived tables in the data foundation. A universe object defined with an analytic function can perform data analysis that would normally require the use of extended syntax at the report level. You might also see improved query performance because calculations are done on the server.

Related Information

[Analytic functions: syntax and examples](#) [page 235]

The generic syntax and examples of analytic functions are given to help you understand how analytic functions can be used.

[Analytic functions: rules, restrictions, and best practices](#) [page 238]

Rules, restrictions, and best practices for using analytic functions are given to help you design the universe.

[Using analytical functions in a business layer object definition](#) [page 239]

To use analytic functions in the business layer, you define the analytic function in the SELECT statement for a measure or dimension.

[Using analytical functions in a derived table definition](#) [page 240]

To use analytic functions in the data foundation, you define the analytic function in the SELECT statement for a derived table.

12.8.1 Analytic functions: syntax and examples

The generic syntax and examples of analytic functions are given to help you understand how analytic functions can be used.

The exact syntax of analytic functions varies depending on the database. Many analytic functions have the following syntax:

Function (arguments) OVER ([<PARTITION BY clause>] [<ORDER BY clause>] [<ROW or RANGE clause>])

Part of analytic function statement	Description
Function (arguments)	The name and arguments of the function that define the calculation.
OVER (OVER signals that this is an analytic function. The OVER clause defines the data over which you want to do the calculation. It has three optional clauses.
PARTITION BY clause	The grouping over which the calculation is applied.
ORDER BY clause	The order of the results to be used in the calculation.

Part of analytic function statement	Description
ROW or RANGE clause)	The interval of records used for calculation.

The PARTITION BY clause lets you define the groups of data over which the function will be calculated. For example:

```
SELECT employee_id, department, COUNT(employee_id) OVER (PARTITION BY department) FROM employee_table
```

This query returns, for each employee, the employee's department and the count of the number of employees in each department. The count is returned with every row (employee) in the result set.

employee_id	department	count
1	Marketing	2
2	Marketing	2
3	Sales	3
4	Sales	3
5	Sales	3

The ORDER BY clause lets you define the order in which the rows are used when applying the calculation. For example:

```
SELECT employee_id, salary, RANK () OVER (ORDER BY salary)
```

This query returns, for each employee, the employee's salary and the employee's overall rank by salary.

employee_id	salary	rank
3	3000	1
2	5000	2
5	6000	3
4	7000	4
1	7200	5

The ROW or RANGE clause lets you define a window or interval of ordered rows to take into account when calculating the function on a given row. For example:

```
SELECT employee_id, salary, SUM(salary) OVER (ORDER BY salary ROWS between unbounded preceding and current row)
```

This query returns, for each employee, the employee's salary and the sum of salaries starting from the lowest salary up to and including the current employee's salary. The results are ordered by salary. The sum for the last row represents the sum of salaries for all employees.

employee_id	salary	sum
3	3000	3000
2	5000	8000
5	6000	14000

employee_id	salary	sum
4	7000	21000
1	7200	28200

The following example uses both the PARTITION BY and ORDER BY clauses:

```
SELECT employee_id, department, salary, RANK() OVER (PARTITION BY department ORDER BY salary)
```

This query returns, for each employee, the employee's department, salary, and rank within the department, ordered by salary within the department.

employee_id	department	salary	rank
2	Marketing	5000	1
1	Marketing	7200	2
3	Sales	3000	1
5	Sales	5000	2
4	Sales	7000	3

The following example uses all three clauses:

```
SELECT employee_id, department, salary, SUM (salary) OVER (PARTITION BY department ORDER BY salary
ROWS between unbounded preceding and current row)
```

This query returns, for each employee, the employee's department and salary, and the sum of salaries within the department starting from the lowest salary in the department up to and including the current employee's salary. The rows are ordered by salary within each department.

employee_id	department	salary	sum
2	Marketing	5000	5000
1	Marketing	7200	12200
3	Sales	3000	3000
5	Sales	5000	8000
4	Sales	7000	15000

Related Information

[About analytic functions](#) [page 234]

You use the information design tool to define analytic functions for objects in a universe in order to calculate, for example, rankings, moving sums or averages, and relative calculations.

12.8.2 Analytic functions: rules, restrictions, and best practices

Rules, restrictions, and best practices for using analytic functions are given to help you design the universe.

The following are some rules and restrictions that apply when using analytic functions in the universe:

- Analytic functions are calculated after the joins are applied, and after the WHERE, HAVING, and GROUP BY clauses are applied. So, for example, the following query returns the value 1:
`COUNT (*) OVER () FROM employee_table WHERE employee_id=312`
- Analytic functions cannot be used in universe conditions and sorts. Disable this usage in the Advanced tab in the object properties. You can, however, use analytic functions in conditions in derived tables.
- Analytic functions cannot appear in a GROUP BY clause. The query expects aggregate functions in the GROUP BY clause.
- Some analytic functions will not work in the same query that contains a GROUP BY clause. Before using analytic and aggregate functions in the same query make sure they are compatible.

The following are some best practices to observe when using analytic functions in the universe:

- In the business layer, keep measures on analytic functions separated from aggregated measures (for example, put them in a different folder).
- Label business layer objects as analytic. Put them in a separate folder or business layer view.
- In the description of the business layer object or derived table, specify any restrictions. For example, that the object cannot be used in a query with the GROUP BY clause (with aggregate functions), or that when using the object you should not apply filters to the query.
- If you anticipate a lot of ad-hoc queries on the universe, consider defining the analytic functionality in derived tables in the data foundation. In the business layer, expose only objects which will always work together.

Related Information

[About analytic functions](#) [page 234]

You use the information design tool to define analytic functions for objects in a universe in order to calculate, for example, rankings, moving sums or averages, and relative calculations.

[Setting where objects can be used](#) [page 267]

12.8.3 Using analytical functions in a business layer object definition

To use analytic functions in the business layer, you define the analytic function in the SELECT statement for a measure or dimension.

Prerequisites

Many analytic functions are listed in the **Database Functions** folder in the **Functions** pane of the SQL Expression Editor. If the function you want to use is not listed, you can add it by updating the extended PRM file.

Note

If a function has both an aggregate and analytic version (for example SUM and SUM OVER), you must define the analytic function in the PRM file if it is not already defined.

To add an analytic function, see the procedure to verify and add analytic function support to PRM files in the *Data Access Guide*. You will need to restart the information design tool after updating the PRM file.

See the rules, restrictions, and best practices for analytic functions in the related topic.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object.
3. In the object properties pane, select the **SQL Definition** tab.
4. Click the **SQL Assistant** button next to the **SELECT** statement to use the SQL editor to build the **SELECT** statement.

Open the **Database Functions** folder in the **Functions** pane and select the desired analytic function. For help on syntax and examples, see the related topic.

5. When you have completed building the SELECT and WHERE statements for the object, save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[Analytic functions: rules, restrictions, and best practices](#) [page 238]

Rules, restrictions, and best practices for using analytic functions are given to help you design the universe.

[Analytic functions: syntax and examples](#) [page 235]

The generic syntax and examples of analytic functions are given to help you understand how analytic functions can be used.

[About analytic functions](#) [page 234]

You use the information design tool to define analytic functions for objects in a universe in order to calculate, for example, rankings, moving sums or averages, and relative calculations.

12.8.4 Using analytical functions in a derived table definition

To use analytic functions in the data foundation, you define the analytic function in the SELECT statement for a derived table.

Prerequisites

Many analytic functions are listed in the **Database Functions** folder in the **Functions** pane of the SQL Expression Editor. If the function you want to use is not listed, you can add it by updating the extended PRM file. To do this, see the procedure to verify and add analytic function support to PRM files in the *Data Access Guide*. You will need to restart the information design tool after updating the PRM file.

Note

To use an analytic function in a derived table definition, updating the extended PRM file is optional, not required.

See the rules, restrictions, and best practices for analytic functions in the related topic.

Procedure

1. Open the data foundation in the editor by double-clicking the data foundation name in the Local Projects View.
2. Insert or edit an existing derived table based on the table that contains the object.
3. Edit the expression for the derived table to build the SELECT statement.
Open the **Database Functions** folder in the **Functions** pane and select the desired analytic function. For help on syntax and examples, see the related topic.
4. When you have completed building the SELECT statement for the table, click **OK** and save the data foundation by clicking the **Save** icon in the main tool bar.

Related Information

[Analytic functions: rules, restrictions, and best practices](#) [page 238]

Rules, restrictions, and best practices for using analytic functions are given to help you design the universe.

[Analytic functions: rules, restrictions, and best practices](#) [page 238]

Rules, restrictions, and best practices for using analytic functions are given to help you design the universe.

[Inserting and editing a derived table](#) [page 171]

[About analytic functions](#) [page 234]

You use the information design tool to define analytic functions for objects in a universe in order to calculate, for example, rankings, moving sums or averages, and relative calculations.

12.9 About aggregate awareness

Aggregate awareness is the ability of a relational universe to take advantage of database tables that contain pre-aggregated data (aggregate tables). Setting up aggregate awareness accelerates queries by processing fewer facts and aggregating fewer rows.

If an aggregate aware object is included in a query, at run time the query generator retrieves the data from the table with the highest aggregation level that matches the level of detail in the query.

For example, in a data foundation there is a fact table for sales with detail on the transaction level, and an aggregate table with sales summed by day. If a query asks for sales details, then the transaction table is used. If a query asks for sales per day, then the aggregate table is used. Which table is used is transparent to the user.

Setting up aggregate awareness in the universe has several steps. See the related topic for more information.

Related Information

[Setting up aggregate awareness](#) [page 241]

12.9.1 Setting up aggregate awareness

Context

This topic outlines the steps to set up aggregate awareness in a relational universe. For links to more detailed information on each step, see the Related Topics.

Procedure

1. The first step is done at the database level. The database administrator must define and load the aggregate tables into the database.

The reliability and usefulness of aggregate awareness in a universe depends on the accuracy of the aggregate tables. They must be refreshed at the same time as all the fact tables.

2. Insert the aggregate tables into the data foundation.
3. Define aggregate aware objects. These are objects in the business layer for which you want queries to use the aggregate tables when possible instead of performing aggregation using non-aggregate tables.

In the SQL expression for the object, define the SELECT statement to use the @Aggregate_Aware function:

```
@Aggregate_Aware(sum(<Aggregate table 1>), ... sum(<Aggregate table n>))
```

In the @Aggregate_Aware function, <Aggregate table 1> is the aggregate table with the highest level of aggregation, and <Aggregate table n> is the aggregate table with the lowest level (the detailed fact table).

4. Specify the incompatible objects for each aggregate table in the universe. In the business layer, use the **Set Aggregate Navigation** command.
5. Resolve any loops in the data foundation using contexts if needed.

Related Information

[About aggregate awareness](#) [page 241]

[Inserting tables into the data foundation](#) [page 151]

[Defining the SQL expression for an object](#) [page 264]

[About @Aggregate_Aware](#) [page 423]

[Setting aggregate navigation](#) [page 242]

[Resolving loops](#) [page 177]

12.9.2 Setting aggregate navigation

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. From the information design tool main menu, select **Actions > Set Aggregate Navigation**.

In the [Aggregate Navigation](#) dialog box, you specify which tables contain objects that are not compatible with aggregate tables containing objects optimized for aggregate awareness:

- If the object is at the same level of aggregation or higher, it is compatible with the aggregate table.
- If the object is at a lower level of aggregation, it is incompatible.
- If the object is unrelated to the aggregate table, it is incompatible.

Note

A measure summed by year is at a higher level of aggregation than a measure summed by quarter.

3. Click an aggregate table in the left pane.
4. In the right pane, select all incompatible objects.
5. Repeat the above steps for each aggregate table in the data foundation.

Note

The dialog box also features a **Detect Incompatibility** button that can guide you in the process of specifying incompatible objects. When you click a table and then click this button, objects considered as

incompatible are automatically selected. Consider the incompatible objects proposed by **Detect Incompatibility** as suggestions, not final choices.

6. When all incompatible objects for all the tables are specified Click **OK**.
7. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information


[About aggregate awareness](#) [page 241]

12.10 Working with business layer objects

The following topics describe how to insert, edit, display, and search for business layer objects.

12.10.1 Inserting a folder

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the folder or analysis dimension in which you want to insert the folder.
To insert a folder at the top level, select the top node (business layer name) in the tree.
3. Click the **Insert object** icon  at the top of the **Business Layer** pane and select **Folder**.
4. Edit the folder properties in the [Folder Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name of the folder as it appears in the Query Panel.
Description	An optional description of the folder.
State	<p>The state of the folder, whether it is Active, Hidden, or Deprecated.</p> <div><p>i Note</p><p>If the state is set to Hidden or Deprecated, the state of the objects in the folder are unchanged, however they do not appear in the Query Panel.</p></div>

Property	Description
Content	A list of objects in the folder that lets you define properties that describe what the object is used for in the query (for result, for filter, for sort). You can change the order of the objects in the folder using the up and down arrow keys to the right of the list.
Custom Properties	Optional custom properties and their values.

5. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]


[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Inserting and editing custom properties](#) [page 276]

12.10.2 Inserting and editing dimensions

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing dimension	Select the dimension.
To insert a dimension	<p>Select the folder or analysis dimension in which you want to insert the dimension. To insert a dimension at the top level, select the top node (business layer name) in the tree.</p> <p>Click the Insert object icon  at the top of the Business Layer pane and select Dimension.</p> <div> <p>i Note</p> <p>Because of an MDX limitation, dimensions that are inserted or copied into an OLAP business layer cannot be used in conditions or sorts. The options Object can be used in Condition and Object can be used in Sort are unavailable.</p> </div>

3. Edit the dimension properties in the [Dimension Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, whether it is Active , Hidden , or Deprecated .
Data Type	The data type of the dimension.
SQL Definition or MDX Definition	The SQL or MDX query expression that defines the object.
Tables (relational only)	The tables associated in the query (by a join in the SQL expression) when returning values for the business layer object. To associate extra tables, click the browse button in the Tables field.
Hierarchy (OLAP only)	<p>The hierarchy that is associated with the dimension.</p> <div> i Note A hierarchy must be specified if you want to insert dimension attributes. </div>
Keys tab (relational only)	<p>The database columns used as primary and foreign keys.</p> <p>Keys allow queries to take advantage of indexes on key columns. Defining keys speeds up data retrieval by optimizing the SQL that is generated for the query. For example in a star schema database if you build a query that filters on a value in a dimension table, the filter can be applied directly on the fact table by using the dimension table foreign key. This avoids inefficient joins to dimension tables.</p>
Advanced tab	<p>Properties that include settings for:</p> <ul style="list-style-type: none"> ○ Access Levels ○ Where in query expressions the object can be used. <div> i Note Because of an MDX limitation, dimensions that are inserted or copied into an OLAP business layer cannot be used in conditions or sorts. The options Object can be used in Condition and Object can be used in Sort are unavailable. </div> <ul style="list-style-type: none"> ○ List of values ○ Display options

Property	Description
Source Information	Descriptive fields that apply to objects used by Data Integrator.
Custom Properties	Optional custom properties and their values.

4. To see the SQL query script for the dimension definition, click **Show Script**.
5. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]
[Inserting dimensions directly from the data foundation](#) [page 246]
[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]
[Defining the SQL expression for an object](#) [page 264]
[Defining the MDX expression for an object](#) [page 264]
[Associating extra tables](#) [page 265]
[Defining keys for dimensions and dimension attributes](#) [page 247]
[Setting object access levels](#) [page 267]
[Setting where objects can be used](#) [page 267]
[Associating a list of values with a business object](#) [page 287]
[Defining custom display formats](#) [page 270]
[About source information for business layer objects](#) [page 276]
[Inserting and editing custom properties](#) [page 276]

12.10.3 Inserting dimensions directly from the data foundation

Context

For business layers based on a data foundation, you can drag and drop objects from the data foundation into the business layer.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View. The data foundation on which the business layer is based displays in the data source pane on the lower right side of the editing tab.
2. Select the objects in the data foundation view that you want to insert:

- To select a table, click the table header.
 - To select multiple tables, click the table headers while holding down the **CTRL** key.
 - To select a column, click the column name in the table.
 - To select multiple columns, click the column names while holding down the **CTRL** key.
3. Drag the selection to the **Business Layer** pane and drop the selection into the desired folder in the business layer. To insert the dimensions at the top level, drop the selection into the top node (business layer name) of the tree.
- When you drag and drop a table, a folder is automatically inserted into the business layer. The folder contains a dimension for every column.
- The SQL expression for each dimension is automatically defined.
4. Save the business layer by clicking the **Save** icon in the main tool bar.

Next Steps

If necessary, turn any of the inserted dimensions into measures using the **Turn into Measure with Aggregation Function** command. For more information, see the related topic.

Related Information

[Turning a dimension or attribute into a measure](#) [page 251]

[Inserting and editing dimensions](#) [page 244]

12.10.4 Defining keys for dimensions and dimension attributes

Prerequisites

Defining keys is available for dimensions and dimension attributes built on a data foundation.


Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Select the dimension or dimension attribute in the **Business Layer** pane.
3. In the *Dimension Properties* pane, click the **Keys** tab.
4. Add keys to the table:


Option	Description
Click Add key .	Adds a key row to the table.
Drag a table column from the data foundation display into the keys table.	Adds a key row to the table and a SELECT statement for the selected column.
Click Detect .	Detects existing key columns in the database and inserts the keys into the table.

You can define one primary and multiple foreign keys for an object. The first key added is the primary key.

- To edit the SELECT statement, click the **SELECT** column.

Enter the **SELECT** statement directly and click  to validate the statement, or click the SQL icon to use the SQL editor to build the statement.

- To enter or edit the WHERE statement, click the **WHERE** column.

Enter the **WHERE** statement directly and click  to validate the statement, or click the SQL icon to use the SQL editor to build the statement.

- Click the **Active** column to enable or disable the key.
- Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[Inserting and editing dimensions](#) [page 244]

[About the SQL/MDX Expression Editor](#) [page 358]

12.10.5 Turning an attribute or measure into a dimension

Context

This task applies to relational business layers.

Procedure

- Open the business layer in the editor by double-clicking the business layer name in the Local Projects View. The business layer objects appear in the **Business Layer** pane and the properties in the editing pane on the right.
- Select the attribute or measure in the **Business Layer** pane.
You can select several attributes or several measures. Click the objects while holding down the **CTRL** key. The command applies to all objects selected.

3. Right-click the selection and select **Turn into Dimension**.
Each dimension is created in the folder of the original object.
4. Save the business layer by clicking the **Save** icon in the main tool bar.


Related Information

[Inserting and editing dimensions](#) [page 244]

12.10.6 Inserting and editing measures

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing measure	Select the measure.
To insert a measure	<p>Select the folder or analysis dimension in which you want to insert the measure. To insert a measure at the top level, select the top node (business layer name) in the tree.</p> <p>Click the Insert object icon  at the top of the Business Layer pane and select Measure.</p>

3. Edit the measure properties in the [Measure Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, whether it is Active , Hidden , or Deprecated .
Data Type	The data type of the object.
Projection Function	Defines how further aggregation is done if needed for the report. Projection aggregation is different from, and is done after the aggregation defined for the measure in the SQL or MDX definition. For more information, see the related topic.

Property	Description
SQL Definition or MDX Definition	The SQL or MDX query expression that defines the object.
Tables (relational only)	The tables associated in the query (by a join in the SQL expression) when returning values for the business layer object. To associate extra tables, click the browse button in the Tables field.
MDX Properties (OLAP only)	You can enter values for the following MDX calculation and format properties to be included in the MDX query: <ul style="list-style-type: none"> ○ Solve order ○ Format string ○ Scope Isolation ○ Language
Advanced tab	Properties that include settings for: <ul style="list-style-type: none"> ○ Access Levels ○ Where in query expressions the object can be used ○ List of values ○ Display options
Source Information	Descriptive fields that apply to objects used by Data Integrator.
Custom Properties	Optional custom properties and their values.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[About projection functions](#) [page 251]

[Defining the SQL expression for an object](#) [page 264]

[Defining the MDX expression for an object](#) [page 264]

[Associating extra tables](#) [page 265]

[Setting object access levels](#) [page 267]

[Setting where objects can be used](#) [page 267]

[Associating a list of values with a business object](#) [page 287]

[Defining custom display formats](#) [page 270]

[About source information for business layer objects](#) [page 276]

[Inserting and editing custom properties](#) [page 276]

12.10.6.1 About projection functions

The projection function defines how to re-aggregate a measure locally for a report. Projection functions apply only to SAP BusinessObjects Web Intelligence reports.

Measures can be aggregated at two different times during the query process:

- First, when the query retrieves the data from the data source, the measure is aggregated according to the SQL or MDX definition of the measure.
- Once the data is retrieved, it is possible to change the level of aggregation in the report. For example, the query retrieves Sales aggregated for Country and City. Then in the Web Intelligence report, Sales is reported only by Country. The projection function defines how to do the local aggregation that is needed to project the data onto the report.

If the projection function is **Sum**, the measure will be summed locally in the report. If the projection function is **Delegated**, the projection function requests that the aggregation be done in the database instead of aggregating locally.

Related Information

[Inserting and editing measures](#) [page 249]

12.10.7 Turning a dimension or attribute into a measure

Context

This task applies to relational business layers.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View. The business layer objects appear in the **Business Layer** pane and the properties in the editing pane on the right.
2. Select the dimension or attribute in the **Business Layer** pane.
You can select several dimensions or several attributes. Click the objects while holding down the **CTRL** key. The command applies to all objects selected.
3. Right-click the selection and select **Turn into Measure with Aggregation Function**.
The aggregation functions that are valid for the data type of the dimension or attribute are listed in a submenu.

Note

If you selected multiple objects with different data types, the aggregation functions for all data types are available, but may not be valid for all objects in the selection.

4. Select the aggregation function for the measure, or **None**.

The SELECT statement in the SQL definition is updated to aggregate the values using the selected function. The data type of the resulting measure is automatically changed, if appropriate. For example, if the original object is of type **DateTime**, and you turn this into a measure with the aggregation function **Count**, the resulting measure has a data type of **Numeric**.

The projection function is automatically set according to the aggregation function selected:

Aggregation Function	Projection Function
Sum	Sum
Count	Sum
Max	Max
Min	Min
Average	Delegated
None	Delegated

For more information about projection functions, see the related topic.

You receive a message listing any objects for which the selected aggregation function is invalid. In this case, the object is turned into a measure, but the new aggregation function is ignored and the projection function is set to **Delegated**.

Note

If any dimension you select is included in a navigation path, you receive a warning that the navigation path will be impacted by the change. If you proceed to turn the dimension into a measure, the dimension is automatically removed from the navigation path.

5. Save the business layer by clicking the **Save** icon in the main tool bar.

Next Steps

You can change the aggregation function at any time by editing the SELECT statement for the measure. If necessary, select another projection function directly in the **Projection Function** dropdown list.

Related Information

[Inserting and editing measures](#) [page 249]


[About navigation paths for objects](#) [page 288]

[About projection functions](#) [page 251]

12.10.8 Inserting and editing attributes

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing attribute	Select the attribute.
To insert an attribute	<div>Select the dimension, measure, hierarchy, or level in which you want to insert the attribute.</div> <div>i Note In an OLAP business layer, when inserting an attribute under a measure, select a measure that was inserted manually into the business layer. The attribute is called Formatted Value Attribute.</div> <div>Click the Insert object icon  at the top of the Business Layer pane and select Attribute or Formatted Value Attribute.</div>

3. Edit the attribute properties in the [Attribute Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, whether it is Active , Hidden , or Deprecated .
Data Type	The data type of the object.
SQL Definition or MDX Definition	The SQL or MDX query expression that defines the object.
Tables (relational only)	The tables associated in the query (by a join in the SQL expression) when returning values for the business layer object. To associate extra tables, click the browse button in the Tables field.
Keys tab (relational dimension attributes only)	The database columns used as primary and foreign keys.

Property	Description
	Keys allow queries to take advantage of indexes on key columns. Defining keys speeds up data retrieval by optimizing the SQL that is generated for the query.
Advanced tab	Properties that include settings for: <ul style="list-style-type: none"> ○ Access Levels ○ Where in query expressions the object can be used ○ List of values ○ Display options
Source Information	Descriptive fields that apply to objects used by Data Integrator.
Custom Properties	Optional custom properties and their values.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Defining the SQL expression for an object](#) [page 264]

[Defining the MDX expression for an object](#) [page 264]

[Associating extra tables](#) [page 265]

[Defining keys for dimensions and dimension attributes](#) [page 247]

[Setting object access levels](#) [page 267]

[Setting where objects can be used](#) [page 267]

[Associating a list of values with a business object](#) [page 287]

[Defining custom display formats](#) [page 270]

[About source information for business layer objects](#) [page 276]

[Inserting and editing custom properties](#) [page 276]

12.10.9 Turning a dimension or measure into an attribute

Context

This task applies to relational business layers.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View. The business layer objects appear in the **Business Layer** pane and the properties in the editing pane on the right.
2. Select the dimension or measure in the **Business Layer** pane.
You can select several dimensions or several measures. Click the objects while holding down the **CTRL** key. The command applies to all objects selected.
3. Right-click the selection and select **Turn into Attribute**.
4. In the *Select a Parent Dimension or Measure* dialog box, select the dimension or measure to which the attributes belong.
5. Save the business layer by clicking the **Save** icon in the main tool bar.


Related Information

[Inserting and editing attributes](#) [page 253]

12.10.10 Inserting and editing filters

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing filter	Select the filter.
To insert a filter	Select the folder or analysis dimension in which you want to insert the filter. To insert a filter at the top level, select the top node (business layer name) in the tree. Click the Insert object icon  at the top of the Business Layer pane and select Filter .

3. Edit the filter properties in the *Filter Properties* pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.

Property	Description
State	The state of the object, whether it is Active , Hidden , or Deprecated .
Filter Type	Either Native (relational business layers only), or Business .
SQL Definition	For native filters, the SQL WHERE expression that defines the object.
Tables (relational only)	The tables associated in the query (by a join in the SQL expression) when returning values for the business layer object. To associate extra tables, click the browse button in the Tables field.
Filter Definition	For business filters, click Edit Filter to define the filter based on objects in the business layer. See the related topic about building a business filter.
Properties tab	<p>When the Use filter as mandatory in query option is selected, the filter is applied to every query using any object in either the universe or the folder, depending on the selected scope (Apply on universe or Apply on folder).</p> <p>When the Apply on list of values option is selected, the filter is applied on list-of-value queries.</p> <p>When Use filter as mandatory in query is unselected, the filter is applied only when explicitly added to the query.</p>
Custom Properties	Optional custom properties and their values.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Defining the SQL expression for an object](#) [page 264]

[Associating extra tables](#) [page 265]

[How to build a business filter](#) [page 308]

[Inserting and editing custom properties](#) [page 276]


12.10.11 Inserting and editing analysis dimensions

Context

Analysis dimensions can be inserted into OLAP business layers only.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing analysis dimension	Select the analysis dimension.
To insert an analysis dimension	Select the business layer name or folder in which you want to insert the analysis dimension. Click the Insert object icon  at the top of the Business Layer pane and select Analysis Dimension .

3. Edit the analysis dimension properties in the [Analysis Dimension Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, whether it is Active , Hidden , or Deprecated . <div>i Note If the state is set to Hidden or Deprecated, the state of the objects in the analysis dimension are unchanged, however they do not appear in the Query Panel.</div>
Type	This property is not currently used.
Default hierarchy	The hierarchy that is taken as the default when the entire analysis dimension is added as a result object in the Query Panel.
Key attribute	This property is not currently used.
Custom Properties	Optional custom properties and their values.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Inserting and editing custom properties](#) [page 276]


12.10.12 Inserting and editing hierarchies

Context

Hierarchies can be inserted into OLAP business layers only.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing hierarchy	Select the hierarchy.
To insert a hierarchy	<p>Select the folder or analysis dimension in which you want to insert the hierarchy. To insert a hierarchy at the top level, select the top node (business layer name) in the tree.</p> <p>Click the Insert object icon  at the top of the Business Layer pane and select Hierarchy.</p>

3. Edit the hierarchy properties in the [Hierarchy Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, whether it is Active , Hidden , or Deprecated .

Property	Description
MDX Definition	The MDX query expression that defines the object.
Advanced tab	Properties that include settings for: <ul style="list-style-type: none"> ○ Access Levels ○ Where in query expressions the object can be used ○ List of values ○ Display options
Source Information	Descriptive fields that apply to objects used by Data Integrator.
Custom Properties	Optional custom properties and their values.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Next Steps

The following objects can be added to the hierarchy:

- Levels
- Attributes
- Named sets
- Calculated members

Related Information

[About business layer objects](#) [page 215]

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Defining the MDX expression for an object](#) [page 264]

[Setting object access levels](#) [page 267]

[Setting where objects can be used](#) [page 267]

[Associating a list of values with a business object](#) [page 287]

[Defining custom display formats](#) [page 270]

[About source information for business layer objects](#) [page 276]

[Inserting and editing custom properties](#) [page 276]

[Inserting and editing hierarchy levels](#) [page 260]

[Inserting and editing attributes](#) [page 253]

[Inserting and editing named sets](#) [page 261]

[Inserting and editing calculated members](#) [page 262]


12.10.13 Inserting and editing hierarchy levels

Context

Levels can be inserted into hierarchies in OLAP business layers only.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing level	Select the level.
To insert a level	Select the hierarchy in which you want to insert the level. Click the Insert object icon  at the top of the Business Layer pane and select Level .

3. Edit the level properties in the [Level Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, when it is Active , Hidden , or Deprecated .
Business Type	This property is not currently used.
MDX Definition	The MDX query expression that defines the object.
Advanced tab	Properties that include settings for: <ul style="list-style-type: none">○ Access Levels○ Where in query expressions the object can be used○ List of values○ Display options
Source Information	Descriptive fields that apply to objects used by Data Integrator.
Custom Properties	Optional custom properties and their values.

Related Information

- [About business layer objects](#) [page 215]
- [Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]
- [Defining the MDX expression for an object](#) [page 264]
- [Setting object access levels](#) [page 267]
- [Setting where objects can be used](#) [page 267]
- [Associating a list of values with a business object](#) [page 287]
- [Defining custom display formats](#) [page 270]
- [About source information for business layer objects](#) [page 276]
- [Inserting and editing custom properties](#) [page 276]


12.10.14 Inserting and editing named sets

Context

Named sets can be inserted into OLAP business layers only.

Procedure

- Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
- In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing named set	Select the named set.
To insert a named set	<p>Select the folder, analysis dimension, or hierarchy in which you want to insert the named set. To insert a named set at the top level, select the top node (business layer name) in the tree.</p> <p>Click the Insert object icon  at the top of the Business Layer pane and select Named Set.</p>

- Edit the named set properties in the [Named Set Properties](#) pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.

Property	Description
State	The state of the object, whether it is Active , Hidden , or Deprecated .
Hierarchy	The hierarchy for the named set.
Named Set Type	The named set type: <ul style="list-style-type: none"> ◦ Native named sets are defined with an MDX expression. ◦ Business named sets are defined by selecting members using the Member Selector.
MDX Definition	For native named sets, The MDX query expression that defines the set.
Definition tab	For business named sets, the list of members. To select members: <ol style="list-style-type: none"> 1. Select a hierarchy from the Hierarchy list. 2. Click Edit Members. 3. In the Member Selector, select or unselect members from the given hierarchy to be included or excluded in the named set. For more information on using the Member Selector, see the related topic.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]

[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Defining the MDX expression for an object](#) [page 264]

[About the Member Selector](#) [page 298]


12.10.15 Inserting and editing calculated members

Context

Calculated members can be inserted into OLAP hierarchies only.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, do one of the following:

Option	Command
To edit an existing calculated member	Select the calculated member.
To insert a calculated member	<p>Select the hierarchy in which you want to insert the calculated member.</p> <p>Click the Insert object icon  at the top of the Business Layer pane and select Calculated Member.</p>

3. Edit the calculated member properties in the *Calculated Member Properties* pane. See the related topics for more information about specific properties.

Property	Description
Name	The name (also called caption in OLAP business layers) of the object as it appears in the Query Panel.
Description	An optional description of the object.
State	The state of the object, whether it is Active , Hidden , or Deprecated .
Hierarchy	The hierarchy for the calculated member.
Parent member	The level in the hierarchy under which the calculated member is to appear. If not specified, the member appears at the root level.
Expression	The MDX expression that defines the calculated member. For more information, see the related topic.
MDX Properties	<p>You can enter values for the following MDX calculation and format properties to be included in the MDX query:</p> <ul style="list-style-type: none">○ Solve order○ Format string○ Scope Isolation○ Language

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About business layer objects](#) [page 215]


[Changing the state of an object: Active, Hidden, or Deprecated](#) [page 266]

[Defining the MDX expression for an object](#) [page 264]

12.10.16 Defining the SQL expression for an object


Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object.
3. In the object properties pane, select the **SQL Definition** tab.

4. Enter the **SELECT** statement directly and click  to validate the SELECT statement, or click the **SQL Assistant** button to use the SQL editor to build the statement.

Most measures require an SQL aggregation function to be defined in the SELECT expression, for example:
`sum(efashion."Shop_facts"."Amount_sold").`

At query run time, the aggregation defined for the measure in the SQL is done before the projection aggregation. You define the projection function separately. For more information, see the related topic.

5. Enter the **WHERE** statement directly and click  to validate the SELECT statement, or click the **SQL Assistant** button to use the SQL editor to build the statement.
6. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About the SQL/MDX Expression Editor](#) [page 358]

[About projection functions](#) [page 251]

[Using analytical functions in a business layer object definition](#) [page 239]


To use analytic functions in the business layer, you define the analytic function in the SELECT statement for a measure or dimension.

12.10.17 Defining the MDX expression for an object

Context

You can edit the MDX expression for objects inserted into the business layer. If you want to edit the definition of a native object (an object generated automatically from the cube when the business layer was created), copy the native object and edit the copy.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object.
3. In the object properties pane, select the **MDX Definition** tab.
4. Enter the **Expression** directly and click  to validate the expression, or click the **MDX Assistant** button to use the MDX editor to build the statement.

Note

When inserting a dimension or level, the best practice is to include `.members` in the MDX expression. For example, following is the expression if you insert the dimension Category in the Category hierarchy in the Product analysis dimension:

```
[Product] . [Category] . [Category] .members
```

For some data sources, it is necessary to add `.members` in order to preview members.

5. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[About the SQL/MDX Expression Editor](#) [page 358]

12.10.18 Associating extra tables

Context

The *Associated SQL Tables* dialog box lets you associate tables with an object in the business layer.

All associated tables are included in the query (by a join in the SQL expression) when returning values for the business layer object. For example, if the object is based on City in table City, when you associate Region and Country as extra tables, the City values from Region and Country are included when values are returned for the business object.

Tables that can be associated in the SQL are proposed in the list.

Procedure

1. To include values from an associated table, check the box next to the table name.
2. To stop including values from a table, uncheck the box next to the table name.

Note

The table on which the business object is based is displayed in bold and cannot be unchecked.

Related Information

[Inserting and editing dimensions](#) [page 244]

[Inserting and editing measures](#) [page 249]

[Inserting and editing attributes](#) [page 253]

12.10.19 Changing the state of an object: Active, Hidden, or Deprecated

Context

Objects in the business layer can have three states:

- **Active:** Object is visible in the Query Panel. This is the default state.
- **Hidden:** Object is valid but not available in the Query Panel (used by other objects as a hidden object).
- **Deprecated:** Object is hidden and not valid. One possible use for this state is when the target database field no longer exists, but you want to keep the object for possible future use.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View. The business layer objects appear in the **Business Layer** pane and the properties in the editing pane on the right.
2. Select the object in the **Business Layer** pane.
To select several objects, click while holding down the **CTRL** key.
3. Right-click the selection and select **Change State**.
4. Select the new state.
The new state is applied to all the objects in the selection. If the state is set to **Hidden** or **Deprecated** for a folder or analysis dimension, the state of the objects in the folder are unchanged, however they do not appear in the Query Panel.
5. Save the business layer by clicking the **Save** icon in the main tool bar.

12.10.20 Setting object access levels

Context

The security access level of an object restricts use of the object to users with the appropriate object-access level granted. You can assign the following access levels to an object:

- **Public**
- **Private**
- **Controlled**
- **Restricted**
- **Confidential**

If you assign **Public** then all users can see and use the object. If you assign **Restricted**, then only users that are granted the object-access level of **Restricted** or higher can see and use the object in the Query Panel.

Universe object-access levels are granted to users and groups in the Central Management Console. For more information, see the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object for which you want to set the access level. To select multiple objects, click while holding down the **CTRL** key.
3. Right-click the selection and select **Change Access Level** and select the new access level from the list.
You can also set the access level for an object on the **Advanced** tab of the object properties.
4. Save the business layer by clicking the **Save** icon in the main tool bar.

12.10.21 Setting where objects can be used

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object.
3. In the object properties pane, click the **Advanced** tab.
4. Select or deselect where the object can be used:

Option	Description
Results	When selected, the object can be used in a query.
Conditions	When selected, the object can be used to set in a condition.

Option	Description
	<p>i Note</p> <p>Because of an MDX limitation, this option is not available for dimensions that are inserted or copied in an OLAP business layer.</p>
Sort	<p>When selected, returned values can be sorted.</p> <p>i Note</p> <p>Because of an MDX limitation, this option is not available for dimensions that are inserted or copied in an OLAP business layer.</p>

- Save the business layer by clicking the **Save** icon in the main tool bar.

12.10.22 Setting options for the default list of values

Context

Dimensions, measures, attributes, and hierarchies are associated with a default list of values. You can set options for the default list of values, or associate a custom list of values with the object.

Procedure

- Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
- In the **Business Layer** pane, select the object.
- In the object properties pane, click the **Advanced** tab.
- Select or deselect list of values options:

Option	Description
Force users to filter values before use	<p>If selected, the user running a query using this list of values is required to enter search criteria before getting filtered values for the list of values. Only the values that match the search criteria are returned in the list of values. Characters used to define the matching criteria are:</p> <ul style="list-style-type: none"> * - Matches any number of characters, even zero characters. ? - Matches exactly one character. \ - Escapes the next character allowing you to search for a wildcard character.
Allow users to search values in the database	<p>If selected, the user running a query using this list of values can search for a value in the database. This option is useful when the user performs a search on partial list of values results.</p>

5. To associate a custom list of values, see the related topic. The options defined in the custom list of values apply.
6. Save the business layer by clicking the **Save** icon in the main tool bar.

Related Information

[Associating a list of values with a business object](#) [page 287]

12.10.23 Creating and editing display formats for business layer objects

You can customize display formats for business layer objects with data types of DateTime and Numeric. There are pre-defined formats available to choose from, or you can create your own custom format using the Format Editor.

Note

When you save the business layer, the custom formats you created for objects in that business layer are saved in the Custom category in the Format Editor. The formats are available to other business layers currently open in the information design tool.

When opening a new information design tool session, to make custom formats available to other business layers, open the business layer in which the formats were defined.

Display formats can be created, edited, and deleted for multiple business layer objects at once.

To create a display format

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select one or more objects, right-click the selection and select **Create Display Format**.
If a display format was already created for an object, the available command is **Edit Display Format**.
3. To select a pre-defined format for the objects, select a Format Category (**Date-Time**, **Numeric**, or **Custom**) and select an available format in the list.
4. To define a custom format, follow the task in the related topics.
5. Click **OK** to use the selected format for the objects.
6. Save the business layer by clicking the **Save** icon in the main tool bar.

Deleting display formats

The **Delete Display Format** command is available for objects with a display format already created. Use this command to unselect the display format previously selected for the object.

Note

Deleting the display format in the business layer does not delete a custom format definition. You must delete a custom format from the Format Editor.

Related Information

[Defining custom display formats](#) [page 270]

12.10.23.1 Defining custom display formats

Procedure

1. In the Format Editor, select an available format to use as a basis for the custom format and click **Custom Format**.

A custom format consists of text and tokens. A token is a pre-formatted part of a number or date. For example, **Day: 1-31** is a token that displays the day part of the date as a number between 1 and 31. For more information about tokens, see the related topics.

The Custom Format Editor lists the token categories. Open the categories to see the list of tokens.

2. Drag tokens from the Tokens list into the **Format Definition**.

Tokens appear in the format definition with a rectangle border and a gray background. You can also type text directly into the **Format Definition**.

When defining a numeric format, you can enter a different format to be displayed when the value is negative or equal to zero. If you do not enter a format, the format defined for positive values is used.

3. In the **Undefined** box, you can enter text to display if no value is returned at reporting time. By default, in the case of an undefined value, no text is displayed.
4. To define a display color for a format, click the color choice box at the end of the format definition.
5. To save the custom format, click **OK**.

Related Information

[Date and time format tokens](#) [page 271]

[Number format tokens](#) [page 274]

12.10.23.1.1 Date and time format tokens

Example

Date and time format display

This example shows how the date, Wednesday March 5th 2008, is displayed using different formats defined in the Custom Format Editor.

Format defined with tokens:	Preview display:
[Day name] , [Month name] [Day 01-31] [Year 0000-9999]	Wednesday, March 05 2008
[Month 01-12] / [Day 01-31] / [Year 0000-9999]	03/05/2008
[Capitalized short day name] [Day 01-31] [Capitalized short month name]	Wed 05 Mar
[Day name], week [Week of year 01-53]	Wednesday, week 10
The current date is [Day name], [Month name] [Day 01-31] [Year 0000-9999] . Day name is [Upper case day name] . Month name is [Lower case month name]. The year is [Year 00-99] .	The current date is Wednesday, March 05 2008. Day Name is WEDNESDAY. Month name is march. The year is 08.

List of date and time tokens

Category	Token	Description
Day	Day 01-31	Day in the month with two digits from 01 to 31.
	Day 1-31	Day in the month with one or two digits from 1 to 31.
	Day name	Day name according to the locale, for example, Monday.
	Short day name	Short day name with capitalization according to the locale, for example, Mon.
	Day of year 001-366	Day in the year with three digits from 001 to 366.
	Day of year 01-366	Day in the year with two or three digits from 01 to 366.
	Day of year 1-366	Day in the year with one, two, or three digits from 1 to 366.

Category	Token	Description
	Day of week in month	Day of the week in the month according to the locale, for example, 3 for the 3rd Monday of June.
	Upper-case day name	Day name in upper case, for example, MONDAY.
	Lower-case day name	Day name in lower case, for example, monday.
	Capitalized day name	Capitalized day name, for example, Monday.
	Upper-case short day name	Short day name in upper case, for example, MON.
	Lower-case short day name	Short day name in lower case, for example, mon.
	Capitalized short day name	Capitalized short day name, for example, Mon.
Month	Month 01-12	Month in the year with two digits from 01 to 12.
	Month 1-12	Month in the year with one or two digits from 1 to 12.
	Month name	Month name with capitalization according to the locale, for example, June.
	Short month name	Short month name with capitalization according to the locale, for example, Jun.
	Upper-case month name	Month name in upper case, for example, JUNE.
	Lower-case month name	Month name in lower case, for example, june.
	Capitalized month name	Capitalized month name, for example, June.
	Upper-case short month name	Short month name in upper case, for example JUN.
	Lower-case short month name	Short month name in lower case, for example, jun.
	Capitalized short month name	Capitalized short month name, for example, Jun.
Year and Era	Year 00-99	Year with two digits from 00 to 99.
	Year 0000-9999	Year with four digits from 0000 to 9999.
	Japanese Imperial period and year	Japanese Imperial period and year number, for example, 平成 20 .
	Japanese Imperial period (English) and year	Japanese Imperial period (English abbreviated) and year number, for example, H20.
	Japanese Imperial year number 01-99	Japanese Imperial year number with two digits.
	Japanese Imperial year number 1-99	Japanese Imperial year number with one or two digits.
	Japanese Imperial period	Japanese Imperial period.
	Japanese Imperial year	Deprecated. Returns the same result as Japanese Imperial year number 0-99 token.
	Era	Era abbreviation, for example, AD or BC.

Category	Token	Description
Week	Week of month	Week in the month with one digit from 1 to 6.
	Week of year 01-53	Week in the year (ISO week) with two digits from 01 to 53.
	Week of year 1-53	Week in the year (ISO week) with one or two digits from 1 to 53.
	Year of week of year 0000	ISO year number (consistent with ISO week) with four digits from 0000 to 9999.
	Year of week of year 00	ISO year number (consistent with ISO week) with two digits from 00 to 99.
Quarter and Semester	Quarter number 1-4	Quarter number with one digit from 1 to 4.
	Short quarter name	Quarter short name from Q1 to Q4.
	Quarter name	Quarter name from 1st quarter to 4th quarter.
	Semester 1-2	Semester number from 1 to 2.
Hour	Hour 00-23	Hour in 24-hour format with two digits from 00 to 23.
	Hour 0-23	Hour in 24-hour format with one or two digits from 0 to 23.
	Hour 01-12	Hour in 12-hour format with two digits from 01 to 12.
	Hour 1-12	Hour in 12-hour format with one or two digits from 1 to 12.
	Hour 01-24	Hour in 24-hour format with two digits from 01 to 24.
	Hour 1-24	Hour in 24-hour format with one or two digits from 1 to 24.
	Hour 00-11	Hour in 12-hour format with two digits from 00 to 11.
	Hour 0-11	Hour in 12-hour format with one or two digits from 0 to 11.
Minute	Minutes 00-59	Minutes with two digits from 00 to 59.
	Minutes 0-59	Minutes with one or two digits from 0 to 59.
Second and sub-second	Seconds 00-59	Seconds with two digits from 00 to 59.
	Seconds 0-59	Seconds with one or two digits from 0 to 59.
	Milliseconds 000-999	Milliseconds with three digits from 000 to 999.
	Hundredths of a second 000-999	Hundredths of a second with two digits from 00 to 99.
	Tenths of a second 0-9	Tenths of a second with one digit from 1 to 9.

Category	Token	Description
Time Zone	Time zone	The offset from Coordinated Universal Time, for example, GMT+00:00.
AM/PM	AM/PM	Morning/afternoon abbreviation, capitalized according to locale, for example, AM or PM. Recommended.
	Upper-case AM/PM	Morning/afternoon abbreviation in upper-case, for example, AM or PM.
	Lower-case am/pm	Morning/afternoon abbreviation in lower-case, for example, am or pm.
	Capitalized Am/Pm	Capitalized morning/afternoon abbreviation, for example, Am or Pm. Not recommended.
Separator	Date separator	Deprecated. This token was used as a date separator in Desktop Intelligence and is not recommended. Type the character you wish to use as a date separator directly into the format description, or use a default format.
	Time separator	Deprecated. This token was used as a time separator in Desktop Intelligence and is not recommended. Type the character you wish to use as a time separator directly into the format description, or use a default format.

12.10.23.1.2 Number format tokens

Number format definitions

A number format definition is made of sections:

- the sign (optional)
- the integer value before the decimal separator
- a grouping separator, to be added in the integer value
- the decimal separator (optional)
- the decimal value after the decimal separator (optional)
- the exponential symbol followed by the exponential value (optional)

Two tokens are used to define the number of significant digits to display in the integer, decimal, and exponential values. Each token in the format definition represents a digit to display:

- The mandatory digit token, **0**, displays the digit if it is significant, otherwise displays a zero.
- The optional digit token, **#**, only displays the digit if it is significant.

When determining the significant digits, the integer value and exponential value are evaluated from right to left, and the decimal value is evaluated from left to right. The last **0** or **#** token is mapped to the remaining digits, if any.

Example

Number format display

This example shows how the value -1,234 is displayed using different formats defined in the Format Editor.

Format defined with tokens:	Preview display:
[Sign] [#]	-1234
[Neg. start] [0] [0] [0] [0] [0] [0] [Neg. end]	(001234)
[Sign always] [#] [Dec.Sep.] [0] [0]	-1234.00
[Sign] [#] [Decimal separator] [0] [0] [E+] [0] [0] [0]	-1.23E+003
Revenue: [Sign always] [#] [Decimal separator] [0] [0]	Revenue: -1234.00 €
[Boolean]	true

List of number format tokens

Category	Token	Description
Signs	Sign	Negative sign if the value is negative. Nothing if the value is positive or zero.
	Sign always	Negative sign if the value is negative. Positive sign if the value is positive or zero.
	Negative start	Open parenthesis if the value is negative. Nothing if the value is positive or zero.
	Negative end	Closing parenthesis if the value is negative. Nothing if the value is positive or zero.
Digits	#	Optional digit. Displays digit only if significant.
	0	Mandatory digit. Displays digit if significant, otherwise displays zero.
Separators	Decimal separator	The symbol used to separate the integer and decimal parts of the number. The symbol used is determined by the locale. The decimal separator can be used only once in an expression.
	Grouping	By default, digits are grouped using the rule and the separator defined by the locale. The grouping symbol can be used only once in an expression. It must appear before the decimal separator.
Exponents	E+	Exponent sign in upper case, always signed. Can be used only once in an expression.
	E-	Exponent sign in upper case, signed only if the value is negative. Can be used only once in an expression.
	e+	Exponent sign in lower case, always signed. Can be used only once in an expression.

Category	Token	Description
	e-	Exponent sign in lower case, signed only if the value is negative. Can be used only once in an expression.
Percent	Percent	The value multiplied by 100.
	Percent %	The value multiplied by 100 followed by the percent sign (%). Can be used only once in an expression.
Boolean	Boolean	Localized value of true if the numerical value is not zero; localized value of false if the numerical value is zero.
	True	Always displays the localized value of true.
	False	Always displays the localized value of false.

12.10.24 About source information for business layer objects

The **Source Information** tab in the business layer object properties contains information about universes generated from Data Integrator. Technical descriptions and formulas used to calculate target tables are displayed.

Property	Description
Technical Information	Information about a column, for example the original database name of the concerned column for the object.
Mapping	Initial formula information describing how a column has been specified (used in Data Integrator), for example revenue = column calculated from several sources.
Lineage	Source columns for the formula used to calculate the column in the database.

12.10.25 Inserting and editing custom properties

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object for which you want to insert or edit custom properties
3. In the object properties pane, select the **Custom Properties** tab.
4. To add a custom property, click **Add**.
5. Edit the property object name and value by clicking the column in the list.
6. To delete a property, select it in the list and click **Delete**.
7. Save the business layer by clicking the **Save** icon in the main tool bar.

12.10.26 Showing associated objects

Context

For relational business layers, you can show the objects in the business layer that reference selected data foundation tables and columns.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the data foundation view pane, select the tables or columns for which you want to see the associated business layer objects. To select a table, click the table header. To select a column, click the column name. To select multiple objects, click while holding down the **CTRL** key.
3. Right-click the selection and select **Show Associated Objects**.
All objects that reference the selected data foundation objects are highlighted in the business layer.

Related Information

[About resource dependencies](#) [page 318]

12.10.27 Showing business layer object values

Context

You can show the values in the underlying data source for a business layer object.

Note that for relational business layers:

- If the object references columns in the data foundation for which a filter is defined, the filters are applied.
- You can also show table and column values from the data foundation view in the Business Layer Editor.

The show values command by default opens a tab in the editor to display the values. You can set a preference to have the values open in a dedicated view or a dialog box. For more information, see the related topic.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Right-click an object in the **Business Layer** pane and select **Show Values**.

The show values window appears. To see what you can do in this window, see the related topic on showing values in a data source.

Related Information

[Showing values in a data source](#) [page 185]

[Showing table values](#) [page 184]

[Showing column values](#) [page 186]





[Setting preferences for showing values](#) [page 34]

12.10.28 Searching for business layer objects

Context

The search panel in the **Business Layer** pane displays the results of a search. All object-contextual commands available in the **Business Layer** pane are also available in the search panel.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View. The **Business Layer** browsing pane displays a tree view of the objects in the business layer.
2. Click the **Show/Hide Search Panel** icon  at the top of the **Business Layer** browsing pane. The **Search Objects** panel opens below the tree view of the business layer.
3. Select the object types to include in the search. Click the filter icon  in the **Search Objects** panel. Select the types to include or exclude. The **Search Objects** panel displays only objects of the types selected.
4. To search the list, click the **Show/Hide search bar** icon .
5. In the search text box, enter the text and press the **Enter** key to start the search. The first object that contains the search text is highlighted in the **Search Objects** and **Business Layer** panels. The total number of objects containing the search text displays in the search text box.
6. To highlight the next object found, press the **Enter** key again. Use the **Enter** key to browse through all objects that match the search text.
7. When you have finished searching, click the **Show/Hide Search Panel** icon  again to hide the **Search Objects** panel.

12.11 About business layer views

You can modify the display of business layer objects by using business layer views to restrict the number of objects displayed in the **Business Layer** pane. Use business layer views to group objects that share a business relationship.

Business layer views can be selected in the Query Panel. You can use business layer views to define security to grant or deny the use of business layer objects to certain users or groups. For more information on defining security using business layer views, see the related topic on Business Security Profile Create Query settings.

You can also filter the **Business Layer** pane in the editor by business layer view.

Related Information


[Creating and editing a business layer view](#) [page 279]

[Filtering by business layer view](#) [page 280]

[Business Security Profile Create Query settings](#) [page 345]

12.11.1 Creating and editing a business layer view

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Click the **Manage Business Layer Views** icon  at the top of the **Business Layer** pane. The [Edit Business Layer View](#) dialog box opens.
3. Do one of the following:
 - To add a view, click **New**.
 - To edit an existing view, select the view in the list.

Note

You cannot edit the **Master** view.

4. Edit the view name in the **Name** text box.
5. In the **Objects in view** box, select or clear the check boxes next to objects in the business layer to include or exclude them from the view.

To work with only the objects already included in the view, select **Show selected objects only**.
6. Enter or edit a description for the view in the **Description** text box.
7. Click **OK** to save the changes.

Related Information

[About business layer views](#) [page 279]

12.11.2 Filtering by business layer view

Context

By default all the folders and objects in the business layer are displayed in the **Business Layer** pane of the editor. You can filter what you see in the Business Layer pane using a business layer view.

You must have at least one business layer view defined.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Select the business layer view in the list at the top of the Business Layer pane.
To return to the display of all objects in the business layer, select **Master** from the list.

Related Information

[Creating and editing a business layer view](#) [page 279]

[About business layer views](#) [page 279]

12.12 About parameters

A parameter is a variable in the business layer or data foundation that requires a value at run time. A parameter can have two input types:

- User input as a response to a prompt. The prompt is a question or directive that requires a user to set one or more values to restrict a result set.
- Pre-defined input that specifies a fixed value for the parameter at run time.

Parameters are defined as individual components in a business layer or data foundation, and are available to all objects in the business layer. You use parameter objects in the SQL or MDX definition of an object to prompt a user response or to implement a fixed value response to a query.

i Note

Parameters inserted in the data foundation are inherited by any business layer based on the data foundation. These parameters cannot be edited in the business layer. You must edit them in the data foundation.

The following properties are available for parameters:

Property	Description
Prompt to users	If selected, the user is prompted to enter a value at run time. If cleared, a pre-defined value is entered at runtime for the parameter.
Prompt Text	The text for the prompt question or directive if Prompt to users is selected.
Set values	Available when the Prompt to users option is unselected. Lets you enter one or more values to be used for the parameter at the run time.
Data Type	The data type required for the answer to the prompt.
Allow multiple values	If selected, lets the user select multiple values from the list of values.
Keep last values	If selected, the last value chosen by the user is kept when the prompt is re-run.
Index aware prompt	If selected, the key column is included in the prompt to restrict the values in a list. The key column is not visible to the user.
Associated list of values	A list of values to provide values for the prompt.
Select only from list	If selected, the user is forced to select a member in the list.
Set default value	Lets you select values to be used as default.

Related Information

[Inserting and editing a parameter](#) [page 281]

[Associating a list of values with a prompt defined in the business layer](#) [page 287]

[Reordering objects in the Business Layer Editor](#) [page 291]

12.12.1 Inserting and editing a parameter


Context

The parameter editor can be started from the business layer or data foundation editor tabs.

Note

Parameters inserted in the data foundation are inherited by any business layer based on the data foundation. These parameters cannot be edited in the business layer. You must edit them in the data foundation.

Procedure

1. Click the **Parameters and Lists of Values** tab in the browsing pane of the editor.
2. Do one of the following:
 - To insert a parameter, click the **Insert Parameter** icon  at the top of the **Parameters** pane.
 - To edit a parameter, click the parameter name in the list.

The properties for the parameter appear in the editor to the right of the **Parameters** pane.
3. Edit properties as required. Parameter properties are described in the related topics.

Related Information

[About parameters](#) [page 280]

[Associating a list of values with a prompt defined in the business layer](#) [page 287]

12.12.2 Creating an index-aware prompt

Context

An index-aware prompt takes advantage of indexes on key columns in tables when accessing lists of values. You can define the prompt so that when running the query, you see and select the user-friendly name for the object. When retrieving values, the query uses the key column for better performance.

To create an index-aware prompt, you create a list of values and a parameter in the data foundation or the business layer. For more detailed information on each step, see the related topics.

Procedure

1. Create a list of values based on custom SQL.
 - a) Include both the key column and name column in the SELECT statement, for example:

```
SELECT reservations.Airline_ID, reservations.Airline_Name FROM reservations
```
 - b) In the **Properties** tab in the list-of-values definition, select the row for the name column. Open the drop-down list in **Key Column**. Select the key column.

For example, the **Key Column** for **Airline_Name** is set to **Airline_ID**.

- c) In the row for the key column, select the **Hidden** check box.

For example, **Hidden** is selected for **Airline_ID**.

2. Create a parameter.

- a) In the **Options** tab in the parameter definition, select **Prompt to users** and enter a **Prompt Text**.
- b) In **Associated List of Values**, select the list of values that you created in step 1.
- c) Select the **Select only from list** check box.
- d) Make sure that the **Index aware prompt** check box is selected.

3. If you want to use the parameter in a query filter in the Query Panel, you need to make the corresponding dimension index aware in the business layer by defining a key on the dimension.

Example

Following is an example of how you might use the index-aware prompt in a WHERE clause (for example in the SQL expression for a derived table, calculated column, or an object in the business layer).

```
WHERE reservations.Airline_ID= @Prompt (<parameter name>)
```

Related Information

[Inserting or editing a list of values](#) [page 284]

[Inserting and editing a parameter](#) [page 281]

[Defining keys for dimensions and dimension attributes](#) [page 247]

[About index awareness](#) [page 234]

12.13 About lists of values


A list of values is a list that contains the data values associated with an object. A list of values allows a user to choose values as a response to a prompt when an associated object is included in a query. The list of values allows a data set to be restricted to the selected values.

A list of values is an independent component in the business layer or data foundation and is available to all business objects in the business layer. A list of values can be associated with an object at any time.

Note

Lists of values inserted in the data foundation are inherited by any business layer based on the data foundation. These lists of values cannot be edited in the business layer. You must edit them in the data foundation.

You can define the following types of lists of values:

Type of List of Values	Description
List of values based on business layer objects (available only in the business layer)	<p>The list of values is based on either a query or a custom hierarchy that includes objects in the business layer. The list is based on the values returned by the query or the hierarchy values.</p> <div>  Restriction You can only base a list of values on a custom hierarchy if the data source supports sub-queries. Otherwise, the option is grayed. </div>
Static list of values	The list of values is based on a list of specified values entered manually or imported from a file.
List of values based on custom SQL	The list of values is based on the values returned by a specified SQL expression.

Related Information

[Inserting or editing a list of values](#) [page 284]

[Reordering objects in the Business Layer Editor](#) [page 291]

12.13.1 Inserting or editing a list of values


Context

The list of values editor can be started from the business layer or data foundation editor tabs.

Note


Lists of values inserted in the data foundation are inherited by any business layer based on the data foundation. These lists of values cannot be edited in the business layer. You must edit them in the data foundation.

Procedure

1. Click the **Parameters and Lists of Values** tab in the browsing pane of the editor.
2. Do one of the following:
 - To insert a list of values, click the Insert List of Values icon  at the top of the **Lists of Values** pane and select the type of list of values. The types are described in the related topic about lists of values.
 - To edit a list of values, click the list of values name in the list.

The properties for the list of values appear in the editor to the right of the **Lists of Values** pane.

3. Edit properties and query options as required. The properties vary depending on the type of list of values:

Option	Description
List of values based on business layer objects (available only in the business layer)	<p>To base the list of values on a query:</p> <ol style="list-style-type: none"> 1. In the Definition tab, select List of values based on the query panel. 2. Click Edit Query. 3. In the Query Panel, select objects and define query filters to define the query that returns the list of values required. 4. Click OK. <p>To base the list of values on a custom hierarchy:</p> <ol style="list-style-type: none"> 1. In the Definition tab, select List of values based on a custom hierarchy. 2. Click Add Dimension. 3. Select dimensions from the list to create the hierarchy required for the list of values. The order of dimensions in the list represents the levels in the hierarchy. Use the up and down arrow keys to modify the order. 4. Click OK. <p>To see the values on the defined list, click Preview.</p>
Static list of values	<p>To add values manually:</p> <ol style="list-style-type: none"> 1. In the Definition tab, click Add Column to add columns to the table. Enter the values for the columns in the table. 2. To add rows, click the Add Row icon  on the right side of the table. <p>To populate the list from a file:</p> <ol style="list-style-type: none"> 1. In the Definition tab, click Import. 2. Select a .txt, .csv, .prn, or .asc file to import as values for the static list. 3. Set the Data Separator, Text Delimiter, and Date Format options according to the format of the data in the file. 4. Click OK. <p>You can edit column properties in the Properties tab. For more information on column properties, see the related topic.</p>
List of values based on custom SQL	<ol style="list-style-type: none"> 1. In the Definition tab, click Edit SQL. 2. In the SQL Editor, build an SQL expression to return the required values, and click OK. <p>To see the values on the defined list, click Preview.</p> <p>You can edit column properties in the Properties tab. For more information on column properties, see the related topic.</p>

4. In the **Options** tab, set the query options for the list of values:

Option	Description
Force users to filter values before use	<p>If selected, the user running a query using this list of values is required to enter search criteria before getting filtered values for the list of values. Only the values that match the search criteria are returned in the list of values. Characters used to define the matching criteria are:</p> <ul style="list-style-type: none"> ○ * - Matches any number of characters, even zero characters. ○ ? - Matches exactly one character.

Option	Description
	<ul style="list-style-type: none"> ○ \ - Escapes the next character allowing you to search for a wildcard character.
Allow users to search values in the database	If selected, the user running a query using this list of values can search for a value in the database. This option is useful when the user performs a search on partial list of values results.
Query Execution timeout	If selected, limits the time in seconds that the list of values query runs.
Max number of rows	If selected, you can enter the maximum number of rows to be returned by the list of values query.

5. Save the business layer or data foundation.

Related Information

[About lists of values](#) [page 283]

[List of values column properties](#) [page 286]

12.13.2 List of values column properties

The **Properties** tab in list of values properties lets you edit the column properties on lists of values. You can edit the following properties by clicking the property column in the table of properties:

Property	Description
Column Name	Lets you to edit the column name.
Key Column	Lets you to select a column to be the index-aware key.
Data Type	Lets you select the data type for the column.
Hidden	When selected, the column will not be displayed to the user. For example, you can hide a column that is only used as a key for another column.

Related Information

[About lists of values](#) [page 283]

12.13.3 Associating a list of values with a business object


Context

Associate a list of values with a business object to restrict possible input values when the object is used as a filter in the Query Panel.

By default, the default list of values is associated with an object.

You can associate a custom list of values with the object. The list of values must be available in the business layer (it is on the list in the **Parameters and Lists of Values** tab of the Business Layer Editor).

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. In the **Business Layer** pane, select the object.
3. In the object properties pane, click the **Advanced** tab.
4. Click the business layer object in the **Business Layer** pane.
5. Select the **Associate list of values** option.
6. To associate a custom list of values, click the browse icon  , select the list of values from the list, and click **OK**.
The options defined for the custom list of values override the options for the default list of values.
7. Save the business layer by clicking the **Save** icon in the main tool bar.


Related Information

[Inserting or editing a list of values](#) [page 284]


[Setting options for the default list of values](#) [page 268]

12.13.4 Associating a list of values with a prompt defined in the business layer


Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Click the **Parameters and Lists of Values** tab under the **Business Layer** pane.
3. Click a parameter in the list in the **Parameters** pane, or click the Insert Parameter icon  to define a new parameter.

The properties for the parameter appear in the editor to the right of the **Parameters** pane.

4. Select the **Prompt to users** option.
5. Click the browse button  at the end of the **Associated list of values** field.
6. Select the radio button for type of list of values.

Type	Description
List of values based on a business layer object	Select values for the list of values from an object in the business layer.
List of values defined in the business layer	Select a pre-defined customized list of values. These are the lists of values listed in the Lists of Values pane.

7. Select either the business layer object or a pre-defined list of values and click **OK**.
8. If you want to restrict the values available in the list to default values, select **Set Default Values**, and click the browse button  at the end of the field.
A selection box appears that lists the available values for the selected object or list. Select values on the left to populate the **Selected values** list and click **OK**.

Next Steps

You can now include the prompt and list of values in the SQL or MDX definition of an object in the business layer using the `@Prompt` function with the name of the parameter defined in this procedure: `@Prompt(<parameter name>)`.

Related Information

[Inserting or editing a list of values](#) [page 284]

[About parameters](#) [page 280]

[About lists of values](#) [page 283]

[About @Prompt](#) [page 426]

12.14 About navigation paths for objects

A Navigation path is an object that defines the drill path used in SAP BusinessObjects reporting tools. A drill path is a list of drillable business objects that allow a report analyst to drill down on a dimension.

A navigation path object can be one of two types:

Navigation path type	Description
Default	<p>The path is defined by the hierarchical organization of the business objects in the business layer. If the business layer contains analysis dimensions, the navigation paths include the dimensions under each analysis dimension. Otherwise, the navigation paths are the dimensions under each folder.</p> <p>You can view the default navigation path in the Navigation Paths tab of the business layer editor. The default path cannot be edited.</p>
Custom	You define the path based on the available dimensions.


Related Information

[Inserting a navigation path object into a business layer](#) [page 289]

[Reordering objects in the Business Layer Editor](#) [page 291]

12.14.1 Inserting a navigation path object into a business layer

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Click the **Navigation Paths** tab under the **Business Layer** pane.
3. Select **Custom** at the top of the **Navigation Paths** pane.
4. Click the **Insert Navigation Path** icon .
5. Enter a **Name** and optionally a **Description** for the path.
The name and description are available to display in the query and reporting tools that use the published universe.
6. Click **Add** to select dimensions for the path. Use the up and down arrow keys to change the order of dimensions in the list.
7. Save the business layer.

Related Information

[About navigation paths for objects](#) [page 288]

12.15 About queries in a business layer

A query object is a query that is saved and associated with the business layer. You use the Query Panel to create queries. Queries are cataloged in the **Query** pane of the editor.

Note

Queries can be used within the information design tool to test the business layer and to preview queries. Query objects are not available to reporting and analysis products using the published universe.

Related Information


[Inserting and editing a query in the business layer](#) [page 290]

[Reordering objects in the Business Layer Editor](#) [page 291]

12.15.1 Inserting and editing a query in the business layer

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Click the **Queries** tab under the **Business Layer** pane.
3. In the **Queries** pane, do one of the following:

Option	Command
To edit an existing query	Select the query. Edit the Name and Description in the Query Properties pane. To edit the query results and filters in the Query Panel, click Edit Query .
To insert a query	Click the Insert Query icon  .

4. In the Query Panel, build or edit the query and click **OK**.
The new query is available in the **Queries** pane.

Related Information

[How to build a query](#) [page 297]

12.16 Reordering objects in the Business Layer Editor

Context

You can reorder lists of values, parameters, queries, and custom navigation paths in the business layer editor. The order is saved in the business layer and the universe outline in the Query Panel. For example, when selecting a parameter to drag it to the **Filter Objects** pane, the parameter objects are presented in the custom order. The custom order in the business layer does not affect the order of prompts when querying.

Procedure

1. Open the business layer in the editor by double-clicking the business layer name in the Local Projects View.
2. Click the tab in the **Business Layer** pane depending on the objects you want to reorder: **Parameters and Lists of Values**, **Queries**, or **Navigation Paths**.

To reorder navigation paths, select **Custom**.

3. Drag and drop the objects into the order that you want them to be listed.

For example, in the **Parameters** pane, drag and drop the parameter names into the desired order.

Inherited lists of values and parameters cannot be reordered in the Business Layer Editor. You must reorder them in the Data Foundation Editor. The inherited objects are listed after the business layer objects in their custom order.

4. Save the business layer by clicking the **Save** icon in the main tool bar.

Results

Note

The A-Z sort direction affects the display only in the editor and is lost if you disable sort or close the editor. By contrast, the custom order that you establish by dragging and dropping objects is kept even after you have closed the editor. To restore the custom sort order, click the **Sort direction** icon and select **Disable sort**.

12.17 About refreshing business layers

Refreshing an OLAP business layer

For business layers based on an OLAP cube, the [Refresh Business Layer](#) wizard detects changes in the OLAP cube and applies the changes to the business layer.




On the [Select Options](#) page, you can select which types of changes the wizard should detect in the cube.

Based on the detections, the wizard lists possible update actions in the [Select Actions](#) page. You can select which update actions you want to apply to the business layer.

Before applying the changes, the wizard displays a summary of update actions on the [Refresh Summary](#) page. You can save the summary to a file. You can go back and modify your selection before finishing the wizard.

A summary list is shown with the changes that are proposed in the business layer based on the changes in the cube structure. You can clear and select proposed changes before applying the update.

Note

A refresh can be undone using the undo action. Undo will recover the business layer to its state before the refresh. To undo, from the information design tool main menu, select  **Edit**  **Undo** .

Refreshing an SAP NetWeaver BW multisource-enabled business layer

When objects are added to the InfoProvider of an SAP NetWeaver BW data source, the procedure to update the universe involves several steps. First, you refresh the structure and synchronize tables in the data foundation. You refresh a business layer with new objects from the data source using the **Insert Candidate Objects** command.

Refreshing relational business layers

To update the business layer with changes to the underlying data foundation, you must manually delete and insert objects. To insert objects for new tables, in the Business Layer Editor, you can drag and drop tables from the data foundation view to the business layer browsing pane.

To identify objects based on tables that have been deleted from the data foundation, follow these steps:

1. In the Local Projects View, right-click the business layer and select **Refresh**.
2. Open the business layer in the editor. If an object in the business layer is based on a table that was deleted from or changed in the data foundation, in the **SQL Definition** tab of the object properties, the **Tables** field label is in red and the field contains the message **[Unresolved table]**.

Related Information

[Refreshing an OLAP business layer](#) [page 293]

[Refreshing universes based on SAP Netweaver BW](#) [page 44]

[Inserting candidate objects](#) [page 293]

[Inserting dimensions directly from the data foundation](#) [page 246]

12.17.1 Refreshing an OLAP business layer

Context

Use the [Refresh Business Layer](#) wizard to update a business layer based on changes in the OLAP cube since the business layer was created, or since the last refresh.

Procedure

1. Open the business layer by clicking the business layer name in the Local Project View.
2. From the information design tool main menu, select **► Actions ► Refresh Structure**.
3. Follow the instructions on the wizard pages. For more information about what to do on a particular page, click the help icon.

Related Information

[About refreshing business layers](#) [page 291]

12.17.2 Inserting candidate objects

Prerequisites

Before inserting candidate objects, first run a refresh structure and then synchronize tables in the data foundation.

Context

Inserting candidate objects applies only to business layers that are based on multisource-enabled data foundations on SAP NetWeaver BW connections. The **Insert Candidate Objects** command detects objects that have been added to the data source since the business layer was created, or since the business layer was last updated with the **Insert Candidate Objects** command. The command detects new objects using the SAP NetWeaver BW strategy.

Inserting candidate objects updates the business layer independently of the data foundation. If you do not also refresh the structure and synchronize tables in the data foundation, you can introduce inconsistencies between the data foundation and business layer.

Procedure

1. Open the business layer by double-clicking the business layer name in the Local Projects View.
2. Right-click the business layer name in the Business Layer pane and select **Insert Candidate Objects**.
A list of business layer objects displays. The objects are grouped in folders based on the source tables. Candidate objects are highlighted and pre-selected.
3. Select the objects in the list to insert into the business layer.

You can select objects that exist in the business layer. In this case, the definitions from the data source will overwrite the existing business object definitions.
4. To insert the selected objects into the business layer, click **Finish**, and save the business layer.

The insertion can be undone using the Edit > Undo command on the main menu.

Next Steps

Inserting candidate objects does not detect obsolete objects in the business layer. You must find and delete obsolete objects manually.

Related Information

[About refreshing a data foundation](#) [page 188]

[Synchronizing tables](#) [page 189]

[Refreshing universes based on SAP Netweaver BW](#) [page 44]

12.18 About computing statistics for optimized query execution

For queries on multisource-enabled universes, you can obtain the best performance if accurate table and columns statistics are available for the data federation service. The cost-based optimizer of the data federation service uses these statistics to determine the optimal join method and order.

The **Compute statistics** command optimizes query execution because it allows you to compute and store statistics in the repository for the universe.

You should compute statistics periodically for tables that might change in volume or for which column values are changing frequently.

The following statistics are generated for the optimization process:

- The table row count
- The number of distinct values for the columns

You set the following options:

- Select all tables and columns that were computed before a certain date
- Select all tables and columns that were never computed
- Select every table and column
- Unselect every table and column

Related Information

[Computing statistics for a multisource-enabled universe](#) [page 295]

12.18.1 Computing statistics for a multisource-enabled universe

Context

You can compute statistics only for universes based on a multisource-enabled data foundation.

Procedure

1. Do one of the following:

Option	Command
To compute statistics from the published universe	In the Repository Resources View, open a session on the repository where the universe is published. Right-click the universe and select Compute Statistics .
To compute statistics from the business layer	Open the business layer in the editor by clicking the business layer name in the Local Projects View. Right-click the business layer name in the Business Layer pane and select Compute Statistics .

2. In the [Compute Statistics](#) dialog box, select the tables and columns for which to compute statistics.
When you select a table, all the columns in the table are selected.
3. Click **Compute**.
The statistics are calculated and stored in the repository. For large databases, this process can take several minutes or longer. While the computation is in progress, you can close the window and perform other tasks in the information design tool.

Related Information

[About computing statistics for optimized query execution](#) [page 294]

13 Using the Query Panel

Use the Query Panel to build, test, and preview the results of queries on a business layer or published universe.

In the information design tool, you can start the Query Panel in the following ways:

- To insert a query into the business layer.
- To open an existing query in the business layer.
- To run a query on a universe published in a repository.

Related Information

[Inserting and editing a query in the business layer](#) [page 290]

[Running a query on a universe published in a repository](#) [page 98]

[How to build a query](#) [page 297]

13.1 How to build a query

Prerequisites


This procedure assumes that you have opened the Query Panel on a business layer or published universe. See the related topic on Using the Query Panel.

Context

You can use this procedure to run queries on published universes, but to save the query you must start the Query Panel from the Business Layer Editor **Queries** pane.

For links to more detailed information about each step, see the Related Topics.



Procedure

1. To select the objects you want to include in the query, drag objects from the business layer on the right into the **Result Objects** pane.
2. For hierarchy result objects, select members to include or exclude in the results. To open the Member Selector, click the arrow to the right of the hierarchy object name: .
3. To filter the results of the query, drag objects from the business layer into the **Filter Objects** pane.

If a mandatory filter is defined on an object, the filter is triggered when you add the object to the **Result Objects** pane. The mandatory filter is visible in the query script, but not in the **Filter Objects** pane.

Non-mandatory pre-defined filters are listed in the business layer. You can drag these pre-defined filters into the **Filter Objects** pane to limit the results. The filter is visible in the query script.

You can also build business filters, including filters that use prompts. For detailed information, see the related topics.

4. For relational universes, you can build combined queries. To open the **Combined Queries** pane, click the  icon.
5. To set query properties, click the  icon.
6. To see or edit the query script, click **View Script**.
7. To preview the query results, click the refresh button in the **Data Preview** pane.

You can profile the values in the result columns. In the **Data Preview** pane, click the **Advanced Preview** icon



To change the layout of hierarchical data, click the **Result set display options** icon  and select an option from the list:

Option	Description
Flat layout	Displays repeated values for a level in every row.
Hierarchical layout	Displays repeated values once for a level.

8. To save the query, click **OK**.
The **OK** button is only available when running the Query Panel from the Business Layer Editor. The query is saved in the business layer and can be executed or edited from the **Queries** pane.

Related Information

[Using the Query Panel](#) [page 297]

[About the Member Selector](#) [page 298]

[How to build a business filter](#) [page 308]

[Filtering data using prompts](#) [page 310]

[Setting query properties](#) [page 313]

[Viewing and editing the query script](#) [page 314]

[Profiling column values in the query panel](#) [page 315]

13.2 About the Member Selector

The Member Selector lets you visualize and select members in a hierarchy. Use the Member Selector to:

- Select the members you want to appear in the query result set.
- Define members that will be excluded from queries.
- Define prompts to allow the selection of members to appear in the query each time you run the query.
- Select the members for a named set.
- Select the members when defining a Business Security Profile filter.

You open the Member Selector from hierarchy objects that you include in queries in the Query Panel. The Member Selector opens automatically when you edit named sets or filters for a Business Security Profile on a hierarchical business layer.

Related Information

[Selecting hierarchy members](#) [page 300]

[About selecting hierarchy members](#) [page 299]

13.2.1 About selecting hierarchy members

In the Member Selector, you can select members in several ways:

- Select members explicitly in the hierarchy. For example, explicitly select the [California] and [Los Angeles] members of the [Geography] hierarchy.
- Select members implicitly using hierarchy relationships. For example, to select US states, you can select the child members of the [US] member.
- Select members included in a named set, for example Top Cities by Revenue, to include the cities that generate the most revenue.
- Select all members in a hierarchy level.
- Select all members up to a certain level in the hierarchy.
- Select calculated members.

The Member Selector contains three tabs:

Tab	Description
Members	Displays the members arranged hierarchically. Use this tab to select members explicitly, by hierarchical relationships, and by specifying all members up to a given level.
Metadata	Shows the hierarchy levels (if the hierarchy supports named levels), named sets, and calculated members.
Prompts	Lets you define and modify prompts.


For information on how to select, display, search for, and sort hierarchy members, see the related topics.

Related Information

[Selecting hierarchy members](#) [page 300]
[Selecting members by hierarchy relationship](#) [page 301]
[Selecting hierarchy members by level](#) [page 302]
[Selecting named sets](#) [page 303]
[Selecting calculated members](#) [page 303]
[Searching for hierarchy members](#) [page 304]
[Excluding hierarchy members](#) [page 305]
[Defining a prompt to select members](#) [page 305]
[Showing selected members in the Member Selector](#) [page 306]
[Sorting hierarchy members](#) [page 307]
[Setting display options](#) [page 307]
[Showing estimated child count](#) [page 307]

13.2.2 Opening the Member Selector in the Query Panel

Procedure


1. In the Query Panel, add the hierarchy object to the **Result Objects** pane.
2. To open the Member Selector, click the arrow to the right of the hierarchy object name: .
3. You can now select members in the hierarchy for inclusion or exclusion in a query. For descriptions of different ways to select members, see the related topic.


Related Information

[About selecting hierarchy members](#) [page 299]

13.2.3 Selecting hierarchy members

Procedure

1. In the Member Selector, click the **Members** tab to display the hierarchy members.
2. Select members in the hierarchy display.
3. To select all members in the hierarchy, click the **Select** icon , and select **Select All**.

- To select all members up to a specified level in the hierarchy, click the **Select** icon . You can identify the level in two ways:

Option	Description
Select a named level	This option is only available if the hierarchy has named levels. Select Select All Members until Named Level and select the level from the submenu.
Select a number of levels below the root	Select Select All Members until and select the number of levels from the submenu.

- When you complete your selection, click **OK**.

Results

The selected members appear below the hierarchy object in the **Result Objects** pane of the Query Panel. When you run the query, only those members are included in the query result.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

13.2.4 Selecting members by hierarchy relationship

Procedure

- In the Member Selector, click the **Members** tab to display the hierarchy members.
- In the hierarchy, right-click the member for which you want to define the hierarchy relationship.
- Select the relationship function from the menu:

Note

Children/Descendants and **Parents/Ancestors** are mutually exclusive pairs. You cannot select both the children and the descendants of a member, and you cannot select both the parents and the ascendants of a member.

Relationship Function	Description
Self	Includes only the selected member. This is the default setting.
Children	Includes members one level below the selected member that have the selected member as their parent. The selected member is not included.

Relationship Function	Description
Descendants	Includes all members at all levels below the selected member. The selected member is not included.
Descendants until Named Level...	Includes the members at levels below the selected member until the named level you select. This option is only available if the hierarchy has named levels.
Descendants until...	Includes the members at levels below the selected member until the number of levels you select.
Parent	Includes the member that is one level above the selected member. The selected member is not included.
Ancestors	Includes all members at all levels above the selected member. The selected member is not included.
Siblings	Includes members at the same level that have the same parent as the selected member. The selected member is not included.
Exclude	Excludes members according to the relationship function (Self/Children/Descendants/Parent/Ancestors/Siblings).

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

13.2.5 Selecting hierarchy members by level

Prerequisites

To select members by level, the hierarchy must have named levels.

Procedure

1. In the Member Selector, click the **Metadata** tab to display the hierarchy levels.

Note

If the **Levels** folder does not display in the **Metadata** tab, the hierarchy is not level-based and you cannot select members by level.

2. Select levels in the **Levels** folder.
3. Click **OK**.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

13.2.6 Selecting named sets

Prerequisites

To select members by named set, the hierarchy must have at least one named set defined. Named sets are defined in the business layer of the universe.

Procedure

1. In the Member Selector, click the **Metadata** tab to display the named sets.

Note

If the **Named Sets** folder does not display in the **Metadata** tab, the hierarchy has no named sets defined.

2. Select named sets in the **Named Sets** folder.
3. Click **OK**.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

[Inserting and editing named sets](#) [page 261]

13.2.7 Selecting calculated members

Prerequisites

To select calculated members, the hierarchy must have at least one calculated member defined. Calculated members are defined in the business layer of the universe.

Procedure

1. In the Member Selector, click the **Metadata** tab to display the calculated members.

Note

If the **Calculated Members** folder does not display in the **Metadata** tab, the hierarchy has no calculated members defined.

2. Select calculated members in the **Calculated Members** folder.
3. Click **OK**.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]


[Inserting and editing calculated members](#) [page 262]

13.2.8 Searching for hierarchy members

Context

Use the Search function in the Member Selector to select hierarchy members from a list of search results.

Procedure

1. To open the [Member Search](#) dialog box, in the Member Selector **Members** tab, click the **Search** icon .
2. Enter text to search for in the **Search pattern** box.

You can use wildcards in the search:

Wildcard	Description
*	Matches any string of characters
?	Matches any one character

3. To search for text in the keys, select the **Search Keys** radio button.
4. Click **Search**.
5. To select members from the search results, select the members in the **Search results** table.
6. Click **OK**.

13.2.9 Excluding hierarchy members

Procedure

1. In the Member Selector, select the members that you want to exclude.
You can select members explicitly, by hierarchy relationship, by level, by named set, and calculated members.
The selected members are listed in the **Summary** pane of the Member Selector.
2. In the **Summary** pane, select the **Exclude** option next to the members or member sets you want to exclude.
3. Click **OK**.

Results

Below the hierarchy object in the **Result Objects** pane of the Query Panel, the excluded members appear with a line drawn through the names to indicate that they are excluded from the query.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

[Selecting hierarchy members](#) [page 300]

[Selecting members by hierarchy relationship](#) [page 301]

[Selecting hierarchy members by level](#) [page 302]

[Selecting named sets](#) [page 303]

[Selecting calculated members](#) [page 303]

[Searching for hierarchy members](#) [page 304]

13.2.10 Defining a prompt to select members

Context

You can define a prompt to defer member selection to the time the query is run.

i Note

When selecting members in response to a prompt, you can only select members explicitly. You cannot select members by hierarchy relationship.

Procedure

1. In the Member Selector, click the **Prompt** tab.
2. Select **Enable Parameter** to defer member selection to when the query is run.
You cannot access the other tabs in the Member Selector when the **Enable Parameter** option is selected.
3. Enter text for the prompt in the **Prompt Text** box.
4. If you want the prompt to select the previously-selected values by default when it displays, select **Keep last values selected**.
5. To define default values for the prompt, select **Set default values** and click **Edit**. In the *Select Parameter Values* dialog box, select default values for the prompt and click **OK**.
6. Click **OK**.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

13.2.11 Showing selected members in the Member Selector

Context

In the Member Selector **Members** tab, you can click the **Expand tree to show selections** icon  to show the selected members in the hierarchy display.

The display automatically expands to show the following members:

- Explicitly selected members.
- Members used to select related members. The related members implicitly selected are not necessarily shown. For example, if the member called France was used to select its children, the tree view expands to show France. If the node France contains no explicitly selected members, the node is not expanded the show the implicitly selected children.

➔ Tip

The **Expand tree to show selections** command does not collapse nodes that are already expanded. To reduce the length of the display, close all open nodes in the hierarchy display before clicking the icon.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

13.2.12 Sorting hierarchy members

Context

By default in the Member Selector, hierarchy members are displayed sorted in the order they are stored in the database. To help find members in the hierarchy, you can sort the display in ascending or descending alphabetical order.

In the Member Selector **Members** tab, click the **Sort order** icon  and select the desired sort order.

The members are sorted locally in the Member Selector. Display of members in the query is not affected.

Related Information

[Opening the Member Selector in the Query Panel](#) [page 300]

13.2.13 Setting display options

Context

By default, the Member Selector displays hierarchy member captions. You can set the display options to display unique names, or both captions and unique names.

In the Member Selector **Members** tab, click the **Member display options** icon  and select the desired display option.

13.2.14 Showing estimated child count

Context

The Member Selector estimates the number of children for each member. By default, the estimates are hidden. You can show the estimated child count in the hierarchy display.

In the Member Selector **Members** tab, click the **Show/Hide estimated child count** icon  to toggle the display of child counts.

13.3 Filtering data in the Query Panel

13.3.1 How to build a business filter

Prerequisites

This procedure assumes you are building a business filter either in the Query Panel or the Edit Business Filter dialog box.

Context

Business filters are filters based on objects in the business layer. They limit the data returned in the query.

Query filters have the following structure: filtered object, operator, operand.

```
[Country] InList (US;France)
```

The **[Country]** dimension is the filtered object, InList is the operator, and the list of values (**US;France**) is the operand. The filter removes all values of **[Country]** other than **US** and **France** from the query result.

The following table describes the components of a filter:

Filter Component	Description
Filtered object	The filtered object is the business layer object whose values are filtered. Dimensions, attributes, measures, hierarchies, and hierarchy levels can be used as filtered objects.
Operator	The operator is used to compare the filtered object with the operand. For example, the Equal To operator retains only those values of the filtered object that correspond exactly to the value of the operand.
Operand	The operand supplies the value or values used to filter the filtered object.

Procedure

1. Drag an object from the business layer to the **Filter Objects** pane. This is the filtered object.
2. In the **Filter Objects** pane, select an operator from the list.
3. In the **Filter Objects** pane, select an operand from the list.

Depending on the type of business layer and purpose of the filter, the following types of operands are available:

Operand type	Description
Constant	<p>Use the Constant operand to enter values directly into the filter. For example, you can use a constant to enter France into the filter:</p> <pre>[Country] Equal To France</pre> <p>You can also enter the @Variable function to retrieve the value of a system variable or User Attribute. For example, to filter on the current user login name, enter the constant operand as @Variable('BOUSER'). For more information about @Variable, see the related topic.</p>
List of Values	<p>Use the List of Values operand to select values from the list associated with the filtered object. For example, if the filtered object is [City], you can use the list of values to select one or more of the cities associated with the object.</p>
Object	<p>Use the object operand to specify an object in the business layer. Drag and drop a business layer object to the operand position when defining the filter.</p> <p>i Note</p> <p>You cannot select an object as an operand on some OLAP data sources or if the filtered object is a hierarchy.</p>
Prompt	<p>Use the prompt operand when you want to be prompted for a value when the query is refreshed. See the related topic about filtering using prompts.</p> <p>i Note</p> <p>Prompt operands are not available if the business filter is defined for a Business Security Profile.</p>

- If you want to filter data on more than one criteria, add an additional filter by dragging another object to the **Filter Objects** pane.

By default, the filters are combined with the AND operator. To use the OR operator, double-click the box with the **And** operator.

i Note

The OR operator is not supported for OLAP data sources.

- If you want to nest query filters, drag another business object and drop it onto an existing query filter in the **Filter Objects** pane.

Nesting query filters allows you to create more complex filter conditions than is possible when you combine filters at the same level. When you nest filters, you set the order in which they are evaluated. Nesting filters only makes sense once you have defined two filters at the same level.

Related Information

[Using the Query Panel](#) [page 297]

[Inserting and editing filters](#) [page 255]

[About @Variable](#) [page 430]

[Building a new prompt to filter data](#) [page 311]

13.3.2 Filtering data using prompts

A prompt is a special type of query filter. It is a dynamic filter that displays a question every time you refresh the data in a query. You answer prompts by either typing or selecting the values you want to view before you refresh the data. The query then returns only the values you specified.

Prompts allow multiple users viewing a single document to specify a different sub-set of the database information and display it in the same report tables and charts. Prompts also reduce the time it takes for the data to be retrieved from the database.

When you define a prompt query filter, you can either build a new prompt, or use an existing prompt defined as a parameter in the business layer.

If you define more than one prompt in a query, you can change the order in which prompts are presented. Change prompt order in the query properties.

Merged Prompts

When querying a business layer or universe, similar prompts are merged. For prompts to be merged, the following rules must be true:

- The prompts have the same prompt text.
- The prompts expect answers having the same data type.
- The prompts expect the same number of answers. (The number of answers to be given depends on the operator used to reference the prompt. For example, **Equal To** expects one answer. **Between** expects multiple answers.)

A single prompt message appears for merged prompts. The list of values displayed by the merged prompt is the list associated with the prompt that has the most display property constraints.

Note

All prompts in the query are candidates for merging: parameters defined in the business layer or data foundation, prompts defined as query filters, and prompts defined in the query expression of a business layer object with the @Prompt function.

Related Information

[Building a new prompt to filter data](#) [page 311]

[Using an existing prompt to filter data](#) [page 312]

[Setting query properties](#) [page 313]

[About parameters](#) [page 280]

13.3.2.1 Building a new prompt to filter data

Prerequisites

This procedure assumes you are building a business filter either in the Query Panel or the Edit Business Filter dialog box.

Procedure

1. Drag the object you want to filter with a prompt and drop it onto the **Query Filters** pane.
The query filter appears in outline in the **Query Filters** pane. The outline shows the filtered object, the operator and the type of filter applied to the object. (By default the filter is a constant.)
2. Select the filter operator from the list.

Note

The list of available operators depends on the type of filtered object.

3. Click the arrow at the right of the outline query filter and select **Prompt** from the menu to filter the object using a prompt.
The **Edit Prompt** dialog box appears and the **New Prompt** option is selected by default.
4. Edit the prompt question in the **Prompt Text** box.
5. Select **Prompt with List of Values** to allow the user to select from a list of values when answering the prompt.
The option is only available if the filtered object has an associated list of values in the universe.
6. Select **Select only from list** to restrict the user choice to values from the list of values.
You can select this option only if the **Prompt with List of Values** option is selected.
7. Select **Keep last values** if you want the prompt to propose the last value that the user selected on the previous refresh. The first time the query is run, the default value (if set) is proposed.
8. Select **Optional prompt** to make the prompt optional. If the user does not supply a value for an optional prompt, the prompt is ignored.
9. Select **Set default values** if you want the prompt to propose values by default when it displays.
 - a) To enter or select the default values, click **Edit**.
 - b) If the filter object has an associated list of values, select the default values from the list.

- c) If the filter object has no associated list of values, enter default values.
 - d) Click **OK** to save the default values.
10. Click **OK** to save the new prompt definition.

Related Information

[How to build a business filter](#) [page 308]

[About lists of values](#) [page 283]

13.3.2.2 Using an existing prompt to filter data

Prerequisites

This procedure assumes you are building a business filter either in the Query Panel or the Edit Business Filter dialog box.

Procedure

1. Drag the object on which you want to apply a prompt and drop it onto the **Query Filters** pane.
The query filter appears in outline in the **Query Filters** pane.
2. Select the filter operator from the list.

Note

The list of available operators depends on the type of filtered object.

3. Click the arrow at the right of the Query Filter and select **Prompt** from the menu.
4. In the **Edit Prompt** dialog box, select the **Use Universe Parameter** option.
5. Select an existing parameter.
The list displays only those universe prompts that are compatible with the object you are filtering. For example, the filtered object and the universe prompt must have the same data type.
6. Click **OK** to save the prompt definition.


Related Information

[How to build a business filter](#) [page 308]

[About parameters](#) [page 280]

13.4 Setting query properties

Procedure

1. In the Query Panel, click the **Query properties** toolbar button .
2. Edit the query property settings as required.

Property	Description
Retrieve duplicate rows	When this option is selected, the query returns all related rows, even if there are duplicate rows. If you do not want duplicate rows in the result set, unselect this option.
Retrieve empty rows (only supported in OLAP universes)	<p>An empty row occurs typically in multidimensional queries when the data for the intersection of two or more dimensions does not exist.</p> <p>When this option is selected, the result set includes rows that can contain empty cells.</p> <p>When this option is unselected, the result set contains only rows with non-empty cells.</p>
Max retrieval time	<p>Defines the maximum time (in seconds) that a query can run before the query is stopped. By default, this value is the same as the Limit Execution Time parameter in the universe parameters.</p> <p>When you set this value to 0, this option is disabled.</p> <p>When the Limit Execution Time parameter is lower than this setting, the Limit Execution Time value is used for limiting the query execution time.</p>
Max rows retrieved	<p>Defines the maximum number of rows of data that are displayed when the query runs. The query retrieves all the possible rows, but only displays the first n rows, where n is the maximum number of rows set for this parameter.</p> <p>The administrator can override this setting in the user security profile settings.</p>
Sample result set	This parameter (when supported by the database) samples n database rows, where n is the value set for the sample result set. This method is faster than using the Max rows retrieved parameter.
Reset contexts on refresh	<p>This is only available on relational universes. When this option is selected, when a user refreshes a query that contains contexts, the user must choose the context(s).</p> <p>The user can clear the previously selected contexts by clicking Clear Contexts.</p>

Property	Description
	When this option is not selected, the query is refreshed using the original contexts. If the contexts have been edited since the last run for the query, the user must choose the contexts again since the query is considered as a new query.
Prompt order	When there are several prompts in a query, use this feature to set the order in which prompts are executed in a query. Click a prompt and use the up or down arrow to change the position of the prompt.

- Click **OK** to close the [Query Properties](#) and save the changes.

13.5 Viewing and editing the query script

Context

You can view the query script of a the query you build in the Query Panel. For relational universes, you can also edit the query script.

Procedure

- In the Query Panel, click **View script**.
The query script displays in the [Query Script Viewer](#).
- For OLAP universes, your only option is to click **OK** to close the [Query Script Viewer](#).
- For relational universes, to edit the query script, select the **Use custom query script** option.
 - In the [Query script](#) pane, edit the query.
 - Click **Validate** to check the script syntax.
 - Click **Undo** to undo the last edit to the script.
 - Click **OK** to save and use the edited query script.
The edited query script is used until you unselect the **Use custom query script** option, or close the Query Panel.
- To use the query script generated by the Query Panel, select the **Use the query script generated by your query** option.
- Click **OK** to save the changes.

13.6 Profiling column values in the query panel

Context

You can profile the values for a columns in the query results. Profiling shows graphically (in a pie or bar chart) the number of occurrences of each value of a column. If the column has a filter defined, the filter is applied.

Procedure

1. Open a query in the Query Panel and refresh the results.

2. In the **Data Preview** pane, click the **Advanced Preview** icon .

The Profile Column Values window appears. To see what you can do in this window, see the related topic.

Related Information

[Showing values in a data source](#) [page 185]

[Using the Query Panel](#) [page 297]

14 Checking integrity

14.1 Running check integrity

Context

Use the **Check Integrity** function to verify aspects of the design of your universe or its elements, for example the data foundation, business layer, parameters, and lists of values. You select pre-defined rules that check the validity of the SQL and MDX expressions, as well as for adherence to design restrictions. Running a check integrity helps to avoid problems when running queries and reports on the published universe.

You can set a background integrity check that performs the integrity check automatically whenever you save a resource. For more information, see the related topic about setting preferences for check integrity.

You can run a check integrity at any time for different objects and resources in the information design tool:

- Resources (data foundations, business layers, connections and shortcuts) in the Local Projects View
- Elements in the data foundation and business layer (tables, contexts, business layer objects, queries, parameters, lists of values) in the editor
- Published universes in the Repository Resources View
- Published universes in the Security Editor (to check the validity of security profiles).

Procedure

1. Right-click the resource or object you wish to run a check integrity for and select **Check Integrity**.
2. In the left-hand pane of the *Check Integrity* dialog box, select the rules you want to apply.
3. Click **Check Integrity**.

The results of the check integrity are listed in the right-hand pane of the *Check Integrity* dialog box. The results of a rule check can have one of three severities:

Severity	Description
Error	The check has detected something that will not work. You must resolve the problem.
Warning	Warning about a missing object (for example, a missing key or a missing link).
Information	The check was OK. A green check mark is displayed next to the rule.

Note

You can change the severity of the results of a rule in the information design tool preferences.

4. To save the results in a text file, click **Export**.
5. When you are finished reviewing the results, click **OK**.

Next Steps

After closing the [Check Integrity](#) dialog box, the results of the check integrity can be reviewed in the Check Integrity Problems View until you run the next check integrity. For more information, see the related topic.

Related Information

[Setting preferences for check integrity](#) [page 26]

[Reviewing check integrity problems](#) [page 317]

14.2 Reviewing check integrity problems

Procedure

1. From the information design tool main menu, select **► Window ► Check Integrity Problems ▾**.
The Check Integrity Problems View opens with a list of the results of the latest check integrity.

Note

If check integrity results are available for more than one resource, the **Problems** view displays the results of the resource that is currently active in the editor.

2. To correct a problem, double-click the result in the list.
The editor opens for the object concerned in the result. For example, if the result concerns a problem with table Customer, the Data Foundation Editor opens with the table Customer highlighted.

Results

The list of results remains in the Check Integrity Problems View until you close the view, or run another check integrity.

Related Information

[Running check integrity](#) [page 316]

15 Showing dependencies between resources

15.1 About resource dependencies

For any local resource, you can view its relationship to other local resources, and any dependent universes published in a repository.

Making changes to a resource, such as deleting it from a local project, moving it to another local project, renaming or updating it, may impact other resources that depend on it. You are warned of the impact before you delete or move a resource.

To help you understand the impact of changes and plan your work, commands exist that show the dependencies between resources and their objects.

Dependencies between local resources

The **Show Local Dependencies** command shows the dependencies between resources in a local project.

When you select a resource, two tabs display the dependent and referenced resources: The **Dependent Resources** tab lists resources in the same local project that depend on the selected resource. The **Referenced Resources** tab lists the resources in the same local project that are referenced by the selected resource. See example 1.

The paths to referenced resources are relative, not absolute. That means that if you rename a resource, the referenced resources are assumed to be in the same folder. If the referenced resources are in a different sub-folder, the reference is broken. See example 2.

You can also show local dependencies for any object in the business layer. In this case, for relational business layers, the referenced resources include the data foundation tables and columns that the object is based on.

Example

1: Show local dependencies

The local project **Demo** contains the following resources:

- **Demo_Local_Connection.cnx**
- **Demo_Data_Foundation.dfx**
- **Demo_for_Accounting.blx**
- **Demo_for_Sales.blx**

You want to list all the resources that would be impacted if you change the **Demo_Data_Foundation**. In the Local Projects View, you select the **Show Local Dependencies** command on **Demo_Data_Foundation.dfx**. The following dependencies are displayed:

Dependent Resources	Referenced Resources
\Demo\Demo_for_Accounting.blx	\Demo\Demo_Local_Connection.cnx

Dependent Resources	Referenced Resources
\Demo\Demo_for_Sales.blx	

The two business layers contain references to the data foundation and might contain invalid references if you delete or change **Demo_Data_Foundation**, so they are listed in the **Dependent Resources** tab.

The data foundation references the connection. Any changes to **Demo_Local_Connection** could impact **Demo_Data_Foundation** and its dependent resources.

Now you want to show the dependencies for **Demo_Local_Connection**:

Dependent Resources	Referenced Resources
\Demo\Demo_Data_Foundation.dfx <ul style="list-style-type: none"> • \Demo\Demo_for_Accounting.blx • \Demo\Demo_for_Sales.blx 	

Notice that **Demo_Data_Foundation** and its two dependent business layers are listed as dependent resources. Since a connection is the first resource created when building a universe, no resources reference the connection.

Example

2: Renaming resources

The local project **OLAP_Demo** contains a folder with the following resources:

Folder_One


- **OLAP_Local_Connection.cnx**
- **OLAP_Business_Layer.blx**

Rename **OLAP_Business_Layer.blx** to **OLAP_New_Business_Layer.blx** and show local dependencies. In the Local Projects View, you select the **Show Local Dependencies** command on **OLAP_New_Business_Layer.blx**. The following dependencies are displayed:

Dependent Resources	Referenced Resources
	\OLAP_Demo\Folder_One \OLAP_Local_Connection.cnx

Even though the business layer was renamed, the **OLAP_Local_Connection** is listed as a referenced resource because it is in the same folder.

Now, create **Folder_Two** in the **OLAP_Demo** project and copy **OLAP_New_Business_Layer.blx** to **Folder_Two**. In **Folder_Two**, rename **OLAP_New_Business_Layer.blx** to **OLAP_New2_Business_Layer.blx** and show local dependencies.

Dependent Resources	Referenced Resources
	 \OLAP_Demo\Folder_Two \OLAP_Local_Connection.cnx

The reference to **OLAP_Local_Connection** is broken because when renaming to **OLAP_New2_Business_Layer.blx**, the information design tool assumes referenced resources are in the same folder.

Dependencies between data foundation and business layer objects

When editing a data foundation, you can show local dependencies for any table or column. A list of dependent business layers displays. You can then display a list of the objects per business layer that depend on the selected data foundation object.

When editing a relational business layer, in the data foundation view, you can select tables and columns and show associated objects. This will highlight in the business layer all objects which reference the selected data foundation objects.

Dependencies between local resources and repository resources

The **Show Repository Dependencies** command lists the universes published in a particular repository that are referenced by the selected local resource.

Related Information

[Showing local dependencies](#) [page 320]

[Showing local dependencies in the data foundation](#) [page 187]

[Showing associated objects](#) [page 277]

[Showing repository dependencies](#) [page 321]

15.2 Showing local dependencies

Context

To show resources in the local project that depend on a selected resource:

Procedure

1. In the Local Projects View, select the resource for which you want to show the dependencies.

-
2. Right-click and select **Show Local Dependencies**.

Results

The **Dependent Resources** tab lists the resources in the same local project that contain references to, or depend on, the selected resource.

The **Referenced Resources** tab lists the resources in the same local project that are referenced by the selected resource.

Related Information

[About resource dependencies](#) [page 318]

15.3 Showing repository dependencies

Context

To show universes in a repository that depend on a selected resource:

Procedure

1. In the Local Projects View, select the resource for which you want to show the dependent resources published in a repository.
2. Right-click and select **Show Repository Dependencies**.
3. Select a session for the repository system where the resources are published, and log in.

Results

The published universes in the repository that reference the selected resource are listed.

Related Information

[Opening a session](#) [page 96]

[About resource dependencies](#) [page 318]

16 Publishing resources

16.1 About publishing resources

Publication is the last step in the universe creation process. Using the Publish Universe Wizard, you publish a business layer to either your local file system or a repository.

When you publish a business layer, the wizard exports the business layer and the resources it references (local connection, connection shortcuts, and data foundation), and creates a universe which is then available to users of query, reporting and analysis tools.

Publishing locally

Only business layers built on local connections can be published locally. This can be a business layer based on a local OLAP connection, or a business layer based on a single-source data foundation with a local connection.

The published universe is saved in the local file system folder that you specify.

Publishing to a repository

To secure a universe, you must publish it first to a repository on a Central Management Server (CMS). The universe inherits the object-level security and user security rights defined for the CMS. The data and metadata in the universe are secured by defining security profiles in the information design tool Security Editor.

When you create a connection in a local project, it is an unsecured local connection and must be published before you can publish a business layer that references the connection. To secure a connection, publish it to a repository on a CMS. The Publish Connection Wizard creates the secured connection and provides a connection shortcut for the local project.

To browse and manage resources once they are published to a repository, use the Repository Resources View.

Editing published resources

You cannot edit a published universe directly in the information design tool. To work on it, you must retrieve it using the Retrieve Universe Wizard. The wizard retrieves the universe from the local folder or the repository, separates it into the business layer and the resources it references (local connection, connection shortcuts, data foundation), and creates these resources in a local project where they can be edited.

Connections can only be published to a repository. To edit a published connection, you must edit it from the Repository Resources View.

Related Information

[Publishing a universe](#) [page 324]

[Retrieving a published universe from a repository](#) [page 70]

[Publishing a local connection to the repository](#) [page 326]

[Publishing a local universe to the repository](#) [page 326]

[About universe security](#) [page 328]

[About managing repository resources](#) [page 94]

16.2 Publishing a universe

Prerequisites

To publish a universe to a repository, the business layer must reference one or more secured connection shortcuts. All shortcuts must reference connections defined in the repository where the universe is to be published.

Note

If the business layer references a local connection and you want to publish to a repository, first publish the connection and change the connection reference in the data foundation (relational), or in the business layer (OLAP) to use the connection shortcut. See the related topics for more information.

To publish a universe locally, the business layer must reference only a local connection that is not secured in any repository.

Recommended actions before publishing a universe:

- Save the business layer and all the resources it references.
- If the business layer references resources that are shared, synchronize the project to ensure that all changes are taken into account in the published universe.
- Check integrity of the business layer and, if applicable, the data foundation. The Publish Universe Wizard gives you the option to run a check integrity before publishing.

Procedure

1. Start the Publish Universe Wizard:
 - To publish to a repository, select the business layer in the Local Projects view, right-click the business layer and select **Publish > To a Repository**.
 - To publish to a local folder, select the business layer in the Local Projects view, right-click the business layer and select **Publish > To a Local Folder**.
2. Follow the instructions on the wizard pages. For more information about what to do on a particular page, click the help icon in the lower left corner.

In the Local Projects View, select the business layer and select **Publish > To a Local Folder**.

Results

The universe is created as a .unx file in the local folder or the repository.

Related Information

[Publishing a local connection to the repository](#) [page 326]

[Changing a connection in a data foundation](#) [page 143]

[Changing the data source of a business layer](#) [page 231]

[Running check integrity](#) [page 316]

[Opening a session](#) [page 96]

[About publishing resources](#) [page 323]

16.2.1 Selecting a repository folder

When publishing or retrieving resources on a repository, the wizard displays the folders in the repository in the left-hand pane. The table in the right-hand pane lists the resources in the folder.

When publishing a resource to a repository, navigate to the repository folder in the navigation tree in the left-hand pane. You can insert a folder.

When retrieving a published universe, navigate to the repository folder in the left-hand pane, and select the universe in the universe list in the right-hand pane.

Note

By default, the resources are retrieved into the local project and are secured locally by requiring you to enter the CMS authentication when opening a retrieved data foundation or business layer.

To remove the local security requirement, select the **Save for all users** option.

16.2.2 Selecting a local folder

Context

When publishing or retrieving a resource in a local folder, the wizard prompts you for a local folder.

Procedure

1. Enter the path to a folder accessible from your local machine.
2. To browse the file system and select a folder, click **Browse**.

16.3 Publishing a local connection to the repository

Procedure

1. To start the Publish Connection Wizard, select the connection in the Local Projects View, right-click the connection and select **Publish Connection to a Repository**.
2. Follow the instructions on the wizard pages. For more information about what to do on a particular page, click the help icon in the lower left corner.

Results

The connection is published in the repository. The local connection is deleted from the Local Projects View. You have the choice to create connection shortcut in the local project. To publish a business layer based on this connection, edit the business layer or data foundation to refer to the new shortcut.

Related Information

[Opening a session](#) [page 96]

[About connection shortcuts](#) [page 101]

[Changing a connection in a data foundation](#) [page 143]

[Changing the data source of a business layer](#) [page 231]

[Synchronizing a project](#) [page 89]

16.4 Publishing a local universe to the repository

Context

You can publish a local universe to a repository if you have a secured connection and update the dependent resources to use this connection.

Procedure

1. Create a connection shortcut in the Local Projects View:

Option	Command
Publish the local connection that the universe is based on.	Follow the procedure to publish a local connection. When asked, create a connection shortcut.
Use an existing secured connection to the data source.	Follow the procedure to create a connection shortcut.

2. Change the dependent resource to reference the connection shortcut:

Option	Command
Relational universes	Edit the data foundation and change the connection to use the connection shortcut.
OLAP universes	Edit the business layer and change the connection to use the connection shortcut.

3. In the Local Projects View, right-click the business layer and select ► **Publish** ► **To a Repository** ▾.

Related Information

[Publishing a local connection to the repository](#) [page 326]

[Creating a connection shortcut](#) [page 129]

[Changing a connection in a data foundation](#) [page 143]

[Changing the data source of a business layer](#) [page 231]

17 Managing security

17.1 About universe security

Universe security starts when the universe is published to a repository on a Central Management Server (CMS). Published universes are stored in the Universes folder and secured connections are stored in the Connections folder.

You secure universes based on the users and groups that are defined in the system repository using the Central Management Console (CMC).

As a first level of security, using the CMC, you grant the right to access specific folders, resources, universes, and connections in the repository to specific users and groups. How to define these rights is described in the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

You define another level of security using the information design tool Security Editor. You can restrict the data that is returned in a query using query limits and controls, filters, and row restrictions. You can also grant or deny access to objects and views in the business layer. To create this level of security, you define security profiles for the universe and assign these profiles to users and groups. The basics of how security profiles work are presented in this topic.

Security profiles

A security profile is a group of security settings that apply to a universe published in the repository. The settings control the data that is displayed and modify the parameters defined in the data foundation and/or business layer. Once the profile is assigned to a user or a group, the settings in the profile determine what objects, data, and connections the user sees when connecting to the universe. There are two types of profiles:

- Data Security Profiles have security settings defined on objects in the data foundation and on relational connections.
- Business Security Profiles have security settings defined on objects in the business layer and on OLAP connections.

Multiple profiles can be defined for each universe. The profiles are saved in the repository.

How profiles work

A user of query and reporting tools who has been granted access to a universe using the CMC, and who has no security profiles assigned or inherited, can see all the objects in the universe and all the data returned by those objects.

When you assign a profile to the user, the security settings defined in the profile are applied whenever the user runs a query on the universe.

In the information design tool, security profiles are applied when you run a query from the Repository Resources View or the Security Editor. They are applied according to the user name you used to open the repository session. When you run a query from the Business Layer Editor, security profile settings do not apply.

How multiple profiles are handled

You can assign more than one profile to a user or group. A user might be assigned a profile and also inherit profiles from groups. When more than one profile is assigned to a user, the profiles are aggregated to produce a single group of settings called the net profile.

Aggregation follows priority and restriction levels that you can modify in the Security Editor. You can also see which profiles a user or group inherits, and preview net profiles for a user or group.

Profile maintenance

Profiles are stored independently from the universe itself: changes in the data foundation or business layer of the universe do not affect the profiles when the universe is republished. Similarly, changes in a profile are independent of assignments, so you do not have to reassign a profile when it is modified. It remains assigned, including any changes.

If you republish a universe, run a check integrity on the universe to flag any discrepancies between the universe and its security profiles.

Profiles created for a universe are deleted when the universe is deleted.

Related Information

[Data Security Profile settings](#) [page 337]

[Business Security Profile settings](#) [page 343]

[Security profile aggregation](#) [page 348]

[Displaying profiles assigned to a user and previewing net profiles](#) [page 356]

[Running a query on a universe published in a repository](#) [page 98]

[About the Security Editor](#) [page 332]

17.2 About securing resources in the information design tool

No authentication is required to start the information design tool.

A user can create and edit unsecured resources (data foundations, business layers, connections) in the Local Projects view. The resources are saved in a local project.

Resources are secured when a user shares a local project and its resources, or publishes universes or connections to a repository. Shared projects and published resources are stored securely in a repository on the Central Management Server (CMS).

Application rights are granted in the Central Management Console (CMC). How to define these rights is described in the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

A user with appropriate rights can retrieve a published universe from the repository for editing. Resources can also be retrieved from a shared project during project synchronization. In both cases, the resources are retrieved into the local project and are secured locally by requiring the user to enter the CMS authentication when opening a retrieved data foundation or business layer.

When you start the information design tool, it reopens any resources that were open when you last closed the tool. If secured resources are open, you need to enter your CMS authentication to start the tool.

Note

To remove the local security requirement, you must have the [Save for all Users](#) right granted in the CMC. When a resource is saved for all users, any user can open the resource without entering CMS authentication.

Secured connections cannot be retrieved from the repository and stored locally in the information design tool. Instead, a shortcut to the connection in the repository is stored in the local project. Secured connections must be edited directly in the repository from the Repository Resources View. To be able to access data from a secured connection (for example, show table values or run a query), the user must enter the CMS authentication for the repository where the connection is published. The system uses the authentication to determine what rights the user has for that connection.

Related Information

[About local projects and resources](#) [page 78]

[About connection shortcuts](#) [page 101]

17.3 CMC rights for information design tool users

The application, universe, and connection rights necessary to do tasks in the information design tool are summarized in this topic.

Rights are granted in the Central Management Console (CMC). How to define these rights is described in the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

To open a session on the repository system and do all secured tasks in the information design tool:

- You must have a username and password configured by the system administrator in the CMC for the CMS where the repository is stored.
- You must have the right [Connect the CMS with Information design tool and view this object in the CMC](#) granted in the CMC.

Additional rights needed are listed in the table by task.

Task	Rights required
Publish a connection to a repository	<ul style="list-style-type: none">• Create, modify, or delete connections application right

Task	Rights required
	<ul style="list-style-type: none"> • <i>View objects</i> right on the connections folder • <i>Add objects to the folder</i> right on the connections folder
Edit a secured connection from the Repository Resources View	<ul style="list-style-type: none"> • <i>Create, modify, or delete connections</i> application right • <i>Add objects to the folder</i> on the connection folder (to create) • <i>Edit objects</i> connection right • <i>Download connection locally</i> connection right (relational connections only)
Use a the local middleware driver for a secured connection	<ul style="list-style-type: none"> • <i>Download connection locally</i> connection right (relational connections only)
Publish a universe to a repository	<ul style="list-style-type: none"> • <i>Publish universes</i> application right • <i>View objects</i> right on the universes folder • <i>Add objects to the folder</i> right on the universe folder • <i>Edit objects</i> universe right (to republish)
Retrieve a published universe from a repository	<ul style="list-style-type: none"> • <i>Retrieve universes</i> application right • <i>View objects</i> right on the universe folder • <i>View Objects</i> universe right • <i>Retrieve universe</i> universe right
Edit secured local resources	<ul style="list-style-type: none"> • No rights are required, but the user must provide the CMS authentication of the user who saved the resources.
Unsecure local resources	<ul style="list-style-type: none"> • <i>Save for all users</i> application right • <i>Retrieve universes</i> application right • <i>View objects</i> right on the universes folder • <i>View Objects</i> universe right • <i>Retrieve universe</i> universe right • <i>Save for all users</i> universe right
Open the Security Editor	<ul style="list-style-type: none"> • <i>Administer security profiles</i> application right
Define security profiles	<ul style="list-style-type: none"> • <i>View objects</i> universe right • <i>Edit security profiles</i> universe right
Assign security profiles to users and groups	<ul style="list-style-type: none"> • <i>View objects</i> universe right • <i>Assign security profiles</i> universe right
Run a query on a published universe	<ul style="list-style-type: none"> • <i>View objects</i> universe right

Task	Rights required
	<ul style="list-style-type: none"> • Create and edit queries based on this universe universe right • Data Access universe right • View objects right on underlying connections • Data Access right on underlying connections
Share project resources: <ul style="list-style-type: none"> • Share a local project • Open the Project Synchronization View • Synchronize project resources • Lock and Unlock resources • Rename or delete a shared project 	<ul style="list-style-type: none"> • Share projects application right
Convert a .unv universe stored in the repository	<ul style="list-style-type: none"> • View objects right on the universe folder • Add objects to the folder right on the universe folder • View objects universe right
Compute statistics for a multisource universe	<ul style="list-style-type: none"> • Compute statistics application right • View objects universe right
Delete a universe from the repository	<ul style="list-style-type: none"> • View objects universe right • Delete objects universe right
Delete a connection from the repository	<ul style="list-style-type: none"> • Create, modify, or delete connections application right • View objects connection right • Delete objects connection right

Related Information

[About session management](#) [page 95]




17.4 About the Security Editor

Use the Security Editor to create and edit security profiles, and assign profiles to users and groups. This topic describes how to navigate the Security Editor. For steps to help you build universe security, see [How to secure a universe using security profiles](#) [page 333].

The session name displays on the tab of the Security Editor. If the session name is prefixed by an asterisk, it means you have made changes to the security profiles or assignments in the Security Editor that have not yet been saved in the repository.

The Security Editor can be viewed in two ways: by universe, or by users/groups. Select the tab on the left-hand side of the Security Editor to display the view you want to work with.

- The **Universes / Profiles** tab lets you to do tasks by first selecting a universe in the repository.
- The **Users / Groups** tab lets you to do tasks by first selecting a user or group. The three icons in the **Users / Groups** panel lets you to display users and groups in three ways:

Icon	Description
	Shows only users.
	Shows all groups and users they contain. A group is shown even if it has no groups or users assigned. Groups are displayed as a flat list. This is the default display.
	Shows all groups and the groups and users they contain. Groups are thus displayed with their different parent groups.

Application rights granted to you in the Central Management Console control what tasks you can perform in the Security Editor. For more information, see the Rights Appendix in the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

Related Information

[About universe security](#) [page 328]

[How to secure a universe using security profiles](#) [page 333]

[Changing Security Profile priority](#) [page 341]

[Changing security profile aggregation options](#) [page 355]

[Assigning Security Profiles to users](#) [page 356]

[Displaying profiles assigned to a user and previewing net profiles](#) [page 356]

17.5 How to secure a universe using security profiles

Prerequisites

Before you begin:

- The universe for which you want to create security must be published to a repository. You can create security for .unx universes only.
- Make sure you have the necessary rights defined in the Central Management Console (CMC). See the related topic about CMC rights.

Context

For links to more detailed information on each step, see the Related Topics.

Procedure

1. Open the Security Editor with a session in the repository where the universe is published.
2. Select the universe in the **Universes / Profiles** pane to define security profiles.
3. For relational universes, insert a Data Security Profile to define the following types of security:
 - Define replacement connections to override the connections defined in the universe.
 - Define overrides for the query options and query limits defined in the universe.
 - Restrict data returned to specific rows using a WHERE clause.
 - Define replacement tables.

Note

You can create more than one Data Security Profile for a universe.

4. Insert a Business Security Profile to define the following types of security:
 - Define a replacement connection to override the connection defined in the universe.
 - Restrict objects that appear in the Query Panel to create queries.
 - Restrict objects for which data is returned.
 - Filter data returned in queries.

Note

You can create more than one Business Security Profile for a universe.

5. Save the changes to the security settings in the repository by clicking the save icon in the main tool bar.
6. Select the **Users / Groups** pane to assign the profiles to users and groups.
7. If more than one profile is assigned to a user (either directly or by inheritance), preview the net results of the aggregated profiles.
8. If modifications to the way profiles are aggregated is needed, in the **Universes / Profiles** pane, change profile priority and profile aggregation options.
9. Save the changes to the security settings in the repository by clicking the save icon in the main tool bar.
10. Test the security profiles for a particular user:
 - a) Open the Security Editor using the login information for the user who is being assigned the security profiles.
 - b) In the **Universes / Profiles** pane, right-click the universe and select **Run Query**.

The Query Panel opens. The security profiles assigned to the user are applied.

Note

Because a user needs the "Administer security profiles" application right granted in order to open the Security Editor, this method of testing profiles is limited. The security profile for a user can be tested in a query application, for example Web Intelligence.

Next Steps

If you republish a universe, run a check integrity on the universe to flag any discrepancies between the universe and its security profiles. In the **Universes / Profiles** pane, right-click the universe and select **Check Integrity**.

Related Information

[CMC rights for information design tool users](#) [page 330]

[Opening the Security Editor](#) [page 335]

[Inserting and editing a Data Security Profile](#) [page 336]

[Inserting and editing a Business Security Profile](#) [page 342]

[Assigning Security Profiles to users](#) [page 356]

[Displaying profiles assigned to a user and previewing net profiles](#) [page 356]

[Changing Security Profile priority](#) [page 341]


[Changing security profile aggregation options](#) [page 355]

[How to build a query](#) [page 297]

[Running check integrity](#) [page 316]

17.6 Opening the Security Editor

Procedure

1. On the information design tool toolbar, click the **Security Editor** icon .
2. In the [Open Session](#) dialog box, select the session you want to open.
3. If you are not already logged into the selected session, enter the required information.

Results

The Security Editor opens in a new tab.

Note

You can open more than one session of the Security Editor at one time. The sessions must be on different repositories.

Related Information

[Opening a session](#) [page 96]

[About the Security Editor](#) [page 332]

17.7 Inserting and editing a Data Security Profile

Context

Caution

Changes to security profiles overwrite any previous changes. In the event that more than one user is editing the same universe profiles at the same time, the changes saved last overwrite changes done previously by others.

Procedure

1. In the Security Editor **Universes / Profiles** pane, select the universe.
2. Do one of the following:

Option	Command
To edit an existing profile	Double-click the profile name.
To insert a profile	Right-click the universe name and select Insert Data Security Profile .

3. Define security settings in each of the tabs by clicking the tab you want.

For more information on the Data Security Profile settings, see the related topics.

Note

Clicking the **Reset** button returns the settings on all tabs to the default values as they are defined in the data foundation and business layer.

4. When you have defined all of the settings, click **OK**.
5. To save the changes to the security settings in the repository, click the save icon in the main tool bar.

Related Information

[Opening the Security Editor](#) [page 335]

[Data Security Profile Connections settings](#) [page 338]

[Data Security Profile Controls settings](#) [page 338]

[Data Security Profile SQL settings](#) [page 339]

[Data Security Profile Rows settings](#) [page 340]

[Data Security Profile Tables setting](#) [page 341]

17.7.1 Data Security Profile settings

A Data Security Profile is a group of settings that define security on a published universe using objects in the data foundation and the data connections.

All Data Security Profile settings apply to relational universes only.

Table 1: Security settings in Data Security Profiles

Security Setting	Description
Connections	Defines replacement relational connections.
Controls	Defines replacement query timeout and size limits.
SQL	Defines replacement query options.
Rows	Defines an SQL WHERE clause to restrict rows returned in the query.
Tables	Defines replacement tables.

Each type of Data Security Profile setting is described in a related topic.

Related Information

[Data Security Profile Connections settings](#) [page 338]

[Data Security Profile Controls settings](#) [page 338]

[Data Security Profile SQL settings](#) [page 339]

[Data Security Profile Rows settings](#) [page 340]

[Data Security Profile Tables setting](#) [page 341]

[Security profile aggregation](#) [page 348]

[Inserting and editing a Data Security Profile](#) [page 336]

17.7.2 Data Security Profile Connections settings

Connections settings are defined in the Data Security Profile for relational universes only (multisource-enabled and single-source). Define replacement connections for OLAP universes in the Business Security Profile.

Use the Data Security Profile Connections setting to define replacement connections that can override the connections defined in the universe. Once a user is assigned or inherits a profile that contains a replacement connection, when the user runs a query on the universe, the replacement connection is used instead of the connection defined in the universe.

Only secured connections can be defined as replacement connections. Relational connections fall into one of three types, listed below. The replacement connection must be of the same type as the original connection.

- SAP NetWeaver BW relational databases
- SAS relational databases
- Other relational databases

To define a replacement connection, select the original connection in the table and click **Edit**.

Select a connection in the Connections folder and sub-folders for which you have the [View objects](#) right granted for the repository where you are defining the security profiles.

For multisource universes that rely on multiple connections, you can define a replacement for each connection.

Related Information

[Aggregation of Connections settings](#) [page 349]

17.7.3 Data Security Profile Controls settings

Controls settings can be defined for multisource-enabled and single-source relational universes.

Use the Data Security Profile Controls settings to define replacement query limits to override default limits when retrieving data from the database. Default query limits are set by the universe designer in the business layer. Once a user is assigned or inherits a profile with replacement Controls settings, when the user runs a query, the replacement limits are used instead of the limits defined in the business layer properties.

In the editor for Data Security Profiles, the limits selected and the limit values defined in the business layer are displayed. When you select or unselect a limit, or enter a new value for a limit, the label appearance changes to bold. This shows that the limit is an override and not the default limit defined in the universe.

Query Limit	Possible values
Limit size of result set to	True and a numerical size between 0 and 2147483647 rows False

Query Limit	Possible values
Limit execution time to	True and a numerical size between 0 and 2147483647 minutes False
Warn if cost estimate exceeds	True and a numerical size between 0 and 10000 minutes False

For more information on query limits, see the related topic about business layer properties.

Related Information

[About business layer properties](#) [page 227]

[Aggregation of Controls settings](#) [page 349]

17.7.4 Data Security Profile SQL settings

SQL settings can be defined for multisource-enabled and single-source relational universes.

Use the Data Security Profile SQL settings to define replacement query options. The universe designer defines default query options in the business layer and data foundation properties. Once a user is assigned or inherits a profile with SQL settings, when the user uses the query panel, the replacement options are used instead of the query options defined in the universe.

In the editor for Data Security Profiles, the SQL settings selected in the business layer and data foundation are displayed. When you select or unselect an option, the label appearance changes to bold. This shows that the option is an override and not the default defined for the universe.

Query Option	Possible values
Allow use of subqueries	True False
Allow use of union, intersect and minus operators	True False
Allow complex operands in Query Panel	True False
Multiple SQL statements for each context	True False
Multiple SQL statements for each measure	True False
Allow Cartesian products	True False

For more information on query options, see the related topics about business layer and data foundation properties.

Related Information

[About business layer properties](#) [page 227]

[About data foundation properties](#) [page 180]

[Aggregation of SQL settings](#) [page 350]

17.7.5 Data Security Profile Rows settings

Rows settings can be defined for multisource-enabled and single-source relational universes.

Use Data Security Profile Rows settings to restrict the rows returned in a query. You restrict the rows by defining an SQL WHERE clause for a specified table. Once a user is assigned or inherits a profile with a Rows setting, when the user runs a query on the universe, the defined WHERE clause is added to the SQL generated if the table is referenced in the query.

Note

A user who has the right to edit the generated SQL in the reporting tool can change the WHERE clause generated by the Rows setting. Remember to manage the user's rights in the reporting tool to prevent the user from modifying the SQL.

You can define the WHERE clause for any standard table in the data foundation. The SQL for the WHERE clause can include:

- @Functions such as @Variable and @Prompt
- For multisource enabled universes, references to other tables in any connection defined for the universe
- For multisource-enabled universes, SAP BusinessObjects SQL functions

The SQL for the WHERE clause cannot include:

- Calculated columns
- Derived tables

Related Information

[Aggregation of Rows settings](#) [page 351]

17.7.6 Data Security Profile Tables setting

Tables settings can be defined for multisource-enabled and single-source relational universes.

Use the Data Security Profile Tables setting to define replacement tables. Once a user is assigned or inherits a profile with a Tables setting, when the user runs a query that references the original table, the replacement table is used instead.

The original table can be a standard table or a federated table in the data foundation. The replacement table can be one of the following types of tables:

- Standard table in the data foundation
- Federated table in the data foundation
- Database table in the connection

Alias and derived tables cannot be defined as either the original or the replacement table.

If you want to specify an owner and qualifier for a replacement table in the database, you must enter these in the fields provided. When you specify a replacement table in this way, the table does not need to exist in the database at design time. So for example, you can specify the table in anticipation of a table that exists at query run time. For more information about data foundation table names, see the related topic.

Note

A user who has the right to edit the generated SQL in the reporting tool can change the replacement table name. Remember to manage the user's rights in the reporting tool to prevent the user from modifying the SQL.

Related Information

[About tables in the data foundation](#) [page 148]

[Aggregation of Tables settings](#) [page 352]

17.8 Changing Security Profile priority

Context

Priority is used to aggregate certain security settings if more than one Data Security Profile or Business Security Profile is assigned to a user or group. For more information on profile aggregation, see the related topic.

Procedure

1. In the Security Editor **Universes / Profiles** pane, select the universe.

2. Right-click the universe name and do one of the following:
 - Select **Change Data Security Profile Priority** (this command is only available if the universe has more than one Data Security Profile defined).
 - Select **Change Business Security Profile Priority** (this command is only available if the universe has more than one Business Security Profile defined).
3. In the dialog box listing the Security Profiles, use the arrow buttons to move profiles up and down in the list. The first profile in the list has highest priority.
4. When you have finished prioritizing, click **OK**.
5. To save the changes in the repository, click the save icon in the main tool bar.

Related Information

[Security profile aggregation](#) [page 348]

[Opening the Security Editor](#) [page 335]

17.9 Inserting and editing a Business Security Profile

Context

Caution

Changes to security profiles overwrite any previous changes. In the event that more than one user is editing the same universe profiles at the same time, the changes saved last overwrite changes done previously by others.

Procedure

1. In the Security Editor **Universes / Profiles** pane, select the universe.
2. Do one of the following:

Option	Command
To edit an existing profile	Double-click the profile name.
To insert a profile	Right-click the universe name and select Insert Business Security Profile .

3. Define security settings in each of the tabs by clicking the tab you want.
For more information on the Business Security Profile settings, see the related topics.

Note

Clicking the **Reset** button returns the settings on all tabs to the default values as they are defined in the data foundation and business layer.

4. When you have defined all of the settings, click **OK**.
5. To save the changes to the security settings in the repository, click the save icon in the main tool bar.

Related Information

[Business Security Profile Connections settings](#) [page 344]

[Business Security Profile Create Query settings](#) [page 345]


[Business Security Profile Display Data settings](#) [page 346]

[Business Security Profile Filters settings](#) [page 347]

17.9.1 Business Security Profile settings

A Business Security Profile is a group of settings that define security on a published universe using objects in the business layer.

Table 2: Security settings for Business Security Profiles

Security Setting	Description
Connections	Defines a replacement OLAP connection.
Create Query	Defines the universe views and business layer objects available to the user in the query panel. <div> Note Create Query settings secure metadata only.</div>
Display Data	Grants or denies access to the data retrieved by objects in the business layer when the user runs a query.
Filters	Defines filters using objects in the business layer.

Each type of Business Security Profile setting is described in a related topic.

In the business layer, designers can set the status of objects to **Active**, **Hidden**, or **Deprecated**. When defining profile settings, you have access to all active objects in the business layer. Objects that are hidden or deprecated in the business layer never appear in the query panel or on reports.

Related Information

[Business Security Profile Connections settings](#) [page 344]

[Business Security Profile Create Query settings](#) [page 345]

[Business Security Profile Display Data settings](#) [page 346]

[Business Security Profile Filters settings](#) [page 347]

[Security profile aggregation](#) [page 348]

[Inserting and editing a Business Security Profile](#) [page 342]

17.9.2 Business Security Profile Connections settings

Connections settings are defined in the Business Security Profile for OLAP universes only. Define replacement connections for relational universes in the Data Security Profile.

Use the Business Security Profile Connections setting to define a replacement connection that can override the connection defined in the universe. Once a user is assigned or inherits a profile that contains a replacement connection, when the user runs a query on the universe, the replacement connection is used instead of the connection defined in the universe.

The replacement connection has the following requirements:

- It must be a secured OLAP connection.
- It must refer to the same database type as the original connection (for example MSAS or Essbase).
- It must specify the catalog and the cube in the connection definition.

Restriction

- SAP NetWeaver BW OLAP connections (**BICS Client**) cannot be used as replacement connections.
- A replacement connection cannot use prompted authentication.

When applying the security setting, the catalog and the cube defined in the replacement connection are used.

To define a replacement connection, select the original connection in the table and click **Edit**.

Select a connection in the Connections folder and sub-folders for which you have the [View objects](#) right granted for the repository where you are defining the security profiles.

Related Information

[Aggregation of Connections settings](#) [page 349]

17.9.3 Business Security Profile Create Query settings

Use the Business Security Profile Create Query setting to grant or deny the use of business layer objects in the query panel.

By default, a user with access to the universe granted in the repository can see all universe objects in the query panel. Once the user is assigned or inherits a profile with a Create Query setting, only the views and objects granted by the setting are displayed and can be selected for a query.

If an object is not granted and not denied explicitly, it is denied by default. Unlike objects that are explicitly denied, objects that are denied by default could be granted by inheritance after aggregating Business Security Profiles to determine the net profile for a user. For more information about aggregating profiles, see the related topic.

There are two ways to grant or deny objects:

- By business layer view: Grants or denies all objects in a view. The **All business layer views** option allows you to grant or deny all views defined for the universe.
- By object: You can grant or deny the objects listed below. The **All objects** option allows you to grant or deny all objects in the business layer.
 - Dimensions
 - Attributes
 - Measures
 - Calculated members
 - Filters
 - Prompts
 - Named sets
 - Folders: Grants or denies all objects in the folder.
 - Analysis dimensions: Grants or denies all objects in the dimension.
 - Hierarchies: Grants or denies all objects in the hierarchy.

Note

It is not possible to grant or deny a hierarchy level.

Tip

If most views are allowed, it is easiest to grant all views, and then deny the ones that are not allowed. Using the **All business layer views** and **All objects** options has the advantage that any new view or object defined in the business layer is automatically included in the Create Query setting when the universe is published.

If the **All business layer views** or **All objects** option is used, the settings are aggregated to determine the net setting for this profile, for example:

- If **All business layer views** are denied and one view is granted, this profile denies all views except the one granted.
- If **All business layer views** are granted and one view is denied, this profile grants all views except the one denied.
- If **All objects** are denied and one object is granted, any parent folders in the path to access the object are granted, but only to access the object. The other objects in the parent folders are denied.
- If **All objects** are granted and one object is denied, the parent folders in the path to access the object are denied, but only to prevent access to this object. The other objects in the parent folders are granted.

The objects in a granted view are granted in that view only. If the same object is contained in another view, it is not automatically granted.

Whether or not a user sees a particular object in the query panel is determined after aggregating the Create Query settings in all profiles assigned to the user, and taking into consideration object access level. For more information about aggregating profiles, see the related topic.

Related Information

[Aggregation of Create Query settings](#) [page 352]

17.9.4 Business Security Profile Display Data settings

Use the Business Security Profile Display Data settings to grant or deny access to the data retrieved by objects in the business layer.

By default, a user with access to the universe granted in the repository can see the data retrieved by all universe objects. Once the user is assigned or inherits a profile with a Display Data setting, only the data corresponding to the objects granted by the setting is displayed.

If an object is not granted and not denied explicitly, it is denied by default. Unlike objects that are explicitly denied, objects that are denied by default could be granted by inheritance after aggregating Business Security Profiles to determine the net profile for a user. For more information about aggregating profiles, see the related topic.

The following objects can be granted or denied. The **All objects** option allows you to grant or deny all objects in the business layer.

- Dimensions
- Attributes
- Measures
- Calculated Members
- Named Sets
- Folders: Grants or denies all the objects in the folder.
- Hierarchies

Using the **All objects** option has the advantage that any new object defined in the business layer is automatically included in the display data setting when the universe is published.

If the **All objects** option is used, the settings are aggregated to determine the net setting for this profile, for example:

- If **All objects** are denied and one object is granted, any parent folders in the path to access the object are granted, but only to access the object. The other objects in the parent folders are denied.
- If **All objects** are granted and one object is denied, the parent folders in the path to access the object are denied, but only to prevent access to this object. The other objects in the parent folders are granted.

A user who is denied an object by a Display Data setting might refresh a report containing the denied object. You can specify what the refresh should do in this case by setting the SQL generation parameter `AUTO_UPDATE_QUERY` in the business layer.

- If this parameter is set to No, then refreshing the report generates an error message.
- If this parameter is set to Yes, then the denied objects are removed from the query and from any filters defined in the business layer. Data for other granted objects is retrieved and displayed to the user in a partial report.

Whether or not a user sees data for a particular object is determined after aggregating the Display Data settings in all profiles assigned to the user, and taking into consideration object access level. For more information about aggregating profiles, see the related topic.

Related Information

[Business Security Profile Display Data settings](#) [page 346]

17.9.5 Business Security Profile Filters settings

Use the Business Security Profile Filters setting to define a filter using objects in the business layer, or named member sets. You create and edit filters explicitly for the Business Security Profile using the Security Editor. Filters in the Business Security Profile are not accessible in the business layer. If the Business Security Profile is deleted, the filter or named set is also deleted.

Once the user is assigned or inherits a profile with a Filters setting, the filter is added to the query script (and therefore combined with any filters defined in the business layer) to restrict the data displayed.

Relational universes

For relational universes, you define filters on dimensions and measures in the business layer. You can define compound filters that are linked by the AND and OR operators. You can also define multiple filters to apply to the query.

When a user runs a query, the filters are always applied to the query and to the returned data. This is different from the Data Security Profile Rows setting which only applies if a defined table is referenced in the query.

OLAP universes

For OLAP universes, you define a named set of members. You can include or exclude members for any dimension in the business layer. The excluded members are removed from the query when data is retrieved from the cube.

Note

The filter does not impact the aggregation of values in the report. Only the display of members is filtered.

You can include or exclude members from multiple dimensions. You can also define multiple named sets to apply to the query.

Related Information

[How to build a business filter](#) [page 308]

[About the Member Selector](#) [page 298]

[Aggregation of Filters settings](#) [page 354]

17.10 Security profile aggregation

More than one Data Security Profile or Business Security Profile defined for a universe may be assigned to the same user. Multiple profiles can be directly assigned to a user or group, and be inherited from parent groups. When this happens, the security settings in the different profiles are aggregated to result in one effective Data Security Profile, and one effective Business Security Profile, called net profiles. The settings in the net profiles are applied when the user creates a query or views a report.

Two methods are used for aggregating security settings: priority and restriction level.

Priority is determined by the order that the Security Profiles appear under the universe in the Security Editor. Use the **Change Data Security Profile Priority** and **Change Business Security Profile Priority** commands to set the priority.

Restriction levels (very restrictive, moderately restrictive, less restrictive) define which operators (for example AND, OR) to use to aggregate profiles. You can change these restriction levels in the Security Editor to affect how the profiles are aggregated.

- The less restrictive level is appropriate when security is designed with roles, each role granting new rights to the user.
- The most restrictive level is appropriate when each profile is used to restrict what the user can see.
- The moderately restrictive level uses the most restrictive level for inherited profiles and the less restrictive level for merged profiles.

The rules for inheriting or merging profiles are as follows:

- If the user or group is assigned Profile A and belongs to a group that is assigned Profile B, Profile A and Profile B are inherited.
- If the user or group belongs to a group assigned Profile A and another group assigned Profile B, Profile A and Profile B are merged.
- If the user or group is assigned both Profile A and Profile B, Profile A and Profile B are merged.

The method and operators used to aggregate profile settings vary for the different settings. For detailed information about aggregation for each type of setting, see the related topic.

The Data Security Profile Rows setting and Business Security Profile Filters setting both generate a WHERE clause to filter the query. The Rows setting is applied first. The WHERE clause in the Filters setting is then applied to the results of the first query. In effect, the two WHERE clauses are aggregated with the AND operator.

Related Information

[Aggregation of Connections settings](#) [page 349]
[Aggregation of Controls settings](#) [page 349]
[Aggregation of SQL settings](#) [page 350]
[Aggregation of Rows settings](#) [page 351]
[Aggregation of Tables settings](#) [page 352]
[Aggregation of Create Query settings](#) [page 352]
[Aggregation of Display Data settings](#) [page 353]
[Aggregation of Filters settings](#) [page 354]
[Changing Security Profile priority](#) [page 341]
[Changing security profile aggregation options](#) [page 355]

17.10.1 Aggregation of Connections settings

If more than one Security Profile for a universe is assigned to or inherited by the same user, the connection defined in the Security Profile with the highest priority is used.

i Note

Connections settings for relational universes are in Data Security Profiles and for OLAP universes, in Business Security Profiles. Therefore Data and Business Security Profiles are never prioritized together.

If the relational universe has multiple connections, aggregation of the Connections setting is done for each connection independently.

Related Information

[Changing Security Profile priority](#) [page 341]
[Business Security Profile Connections settings](#) [page 344]
[Data Security Profile Connections settings](#) [page 338]

17.10.2 Aggregation of Controls settings

If more than one Data Security Profile for a universe is assigned to or inherited by the same user, the following rules are used to aggregate the Controls settings. The rules are applied to each query limit to determine the value to be used when the user runs a query or report.

Restriction level	Aggregation Rule
Very Restrictive	<p>The limit is active only if it is selected in all merged and inherited profiles.</p> <p>The value used is the minimum value for the limit among all the merged and inherited profiles.</p>
Moderately Restrictive	<p>The limit is active only if it is selected in all inherited profiles and selected in at least one merged profile.</p> <p>First, the minimum value is determined for the limit comparing the inherited profiles. This value is compared to the values among the merged profiles. The value used is the maximum among these values.</p>
Less Restrictive	<p>The limit is active if it is selected in any merged or inherited profile.</p> <p>The value used is the maximum value for the limit among all the merged and inherited profiles.</p>
Priority (default)	The activation and value of the limit in the Data Security Profile with the highest priority is used.

i Note

For a definition of inherited and merged profiles, see the related topic on security profile aggregation.

Related Information

[Security profile aggregation](#) [page 348]

[Changing security profile aggregation options](#) [page 355]

[Changing Security Profile priority](#) [page 341]

[Data Security Profile Controls settings](#) [page 338]

17.10.3 Aggregation of SQL settings

If more than one Data Security Profile for a universe is assigned to or inherited by the same user, the following rules are used to aggregate the SQL settings. The rules are applied to each query option to determine the value to be used when the user creates a query.

Restriction level	Aggregation Rule
Very Restrictive	The option is active only if it is selected in all merged and inherited profiles.

Restriction level	Aggregation Rule
Moderately Restrictive	The option is active if it is selected in all inherited profiles and selected in at least one assigned profile.
Less Restrictive	The option is active if it is selected in any merged or inherited profile.
Priority (default)	The activation and value of the option in the Data Security Profile with the highest priority is used.

i Note

For a definition of inherited and merged profiles, see the related topic on security profile aggregation.

Related Information

[Security profile aggregation](#) [page 348]

[Changing security profile aggregation options](#) [page 355]

[Changing Security Profile priority](#) [page 341]

[Data Security Profile SQL settings](#) [page 339]

17.10.4 Aggregation of Rows settings

If more than one Data Security Profile for a universe is assigned to or inherited by the same user, the following rules are used to aggregate the Rows settings and determine the WHERE clause to be used when the user runs a query or report.

First, the WHERE clauses for each table are aggregated according to the restriction level:

Restriction level	Aggregation Rule
Very Restrictive (default)	The WHERE clauses in all profiles that apply to the same table are combined with the AND operator.
Moderately Restrictive	Inherited WHERE clauses are aggregated using the AND operator. Merged WHERE clauses are aggregated using the OR operator.
Less Restrictive	The WHERE clauses in all profiles that apply to the same table are combined with the OR operator.

After aggregation according to restriction level, the WHERE clauses for each table are aggregated together with the AND operator to produce the final WHERE clause that is applied to the query.

Note

For a definition of inherited and merged profiles, see the related topic on security profile aggregation.

Related Information

[Security profile aggregation](#) [page 348]

[Changing security profile aggregation options](#) [page 355]

[Data Security Profile Rows settings](#) [page 340]

17.10.5 Aggregation of Tables settings

If more than one Data Security Profile for a universe is assigned to or inherited by the same user, the replacement table defined in the Data Security Profile with the highest priority is used. If settings are defined for multiple tables, the aggregation is done for each table independently.

Related Information

[Changing Security Profile priority](#) [page 341]

[Data Security Profile Tables setting](#) [page 341]

17.10.6 Aggregation of Create Query settings

If more than one Business Security Profile for a universe is assigned to or inherited by the same user, the Create Query settings are aggregated. Object access levels, if they are defined, are applied to determine whether or not a user sees a particular object in the query panel.

First, the list of views that the user can select in the query panel is determined by aggregating the profiles according to the restriction level:

Restriction level	Aggregation Rule
Very Restrictive (default)	The user can select the view in the query panel only if it is granted in all inherited and merged profiles.
Moderately Restrictive	The user can select the view in the query panel only if it is granted in all inherited profiles and granted in at least one merged profile.

Restriction level	Aggregation Rule
Less Restrictive	The user can select the view in the query panel if it is granted in any inherited or merged profile.

Once a view is selected in the query panel, an object appears if it is included in the view, and if it is not explicitly denied after aggregating the profiles according to restriction level:

Restriction level	Aggregation Rule
Very Restrictive (default)	The object is denied if it is explicitly denied in any inherited or merged profile.
Moderately Restrictive	The object is denied if it is explicitly denied in any inherited profile and denied in all merged profiles.
Less Restrictive	The object is denied only if it is explicitly denied in all inherited and merged profiles.

After aggregation, the denied objects are not displayed even if they belong to a granted view. If a folder is denied, then all sub-folders and objects in the folder are denied.

Finally, the access level granted to the user in the Central Management Console determines which of the objects granted by the net Business Security Profile are available in the query panel. The user sees only the objects with an access level lower than or equal to his authorized access level. You assign access levels to objects in the business layer editor.

i Note

For a definition of inherited and merged profiles, see the related topic on security profile aggregation.

For more information about object access levels, see the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

Related Information

[Security profile aggregation](#) [page 348]

[Changing security profile aggregation options](#) [page 355]

[Business Security Profile Create Query settings](#) [page 345]

17.10.7 Aggregation of Display Data settings

If more than one Business Security Profile for a universe is assigned to or inherited by the same user, the Display Data settings are aggregated. Object access levels, if they are defined, are applied to determine whether or not a user sees the data for an object in the business layer.

First, the list of objects that the user can see data for is determined by aggregating the profiles according to restriction level.

Restriction level	Aggregation Rule
Very Restrictive (default)	The data appears only if it is granted in all inherited and merged profiles.
Moderately Restrictive	The data appears only if the object is granted in all inherited profiles and granted in at least one merged profile.
Less Restrictive	The data appears if the object is granted in any inherited or merged profile.

If a folder is denied, then the data for all objects in the folder and its subfolders is denied.

Finally, the access level granted to the user in the Central Management Console determines which of the objects granted by the net Business Security Profile the user sees data for. The user sees data only for the objects with an access level lower than or equal to his authorized access level. You assign access levels to objects in the business layer editor.

Note

For a definition of inherited and merged profiles, see the related topic on security profile aggregation.

For more information about object access levels, see the *SAP BusinessObjects Business Intelligence platform Administrator's Guide*.

Related Information

[Security profile aggregation](#) [page 348]

[Changing security profile aggregation options](#) [page 355]

[Business Security Profile Display Data settings](#) [page 346]

17.10.8 Aggregation of Filters settings

If more than one Business Security Profile for a universe is assigned to or inherited by the same user, the following rules are used to aggregate the filters settings and determine the filter to be added to the query script when the user runs a query or report.

For relational universes, the filters are aggregated according to the restriction level. The resulting filter is added to the WHERE clause applied to the query.

Restriction level	Aggregation Rule
Very Restrictive (default)	The filters in all profiles are combined using the AND operator.
Moderately Restrictive	Inherited filters are aggregated using the AND operator.

Restriction level	Aggregation Rule
	Merged filters are aggregated using the OR operator.
Less Restrictive	The filters in all profiles are combined using the OR operator.

For OLAP universes, the named sets are aggregated according to the restriction level.

Restriction level	Aggregation Rule
Very Restrictive (default)	The user sees a member only if it is included in every named set defined in all profiles.
Moderately Restrictive	The user sees a member if it is included in every named set defined in the inherited profiles, and included in at least one named set defined in the merged profiles.
Less Restrictive	The user sees a member if it is included in any named set defined in any profile.

i Note

For a definition of inherited and merged profiles, see the related topic on security profile aggregation.

Related Information

[Security profile aggregation](#) [page 348]

[Changing security profile aggregation options](#) [page 355]

[Business Security Profile Filters settings](#) [page 347]

17.11 Changing security profile aggregation options

Procedure

1. In the Security Editor **Universes / Profiles** pane, select the universe.
The current aggregation options for the universe display on the lower, right-hand side of the editor.
2. For each security setting, select a new aggregation option from the list.
The options apply only to the universe currently selected.
3. To save the changes in the repository, click the save icon in the main tool bar.

Related Information

[Security profile aggregation](#) [page 348]

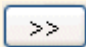
[Opening the Security Editor](#) [page 335]

17.12 Assigning Security Profiles to users

Procedure

1. In the Security Editor **Universes / Profiles** pane, select the universe.
Any currently assigned users or groups appear in the **Assigned Users** list.
2. To assign, select the user or group in the list of users on the right-hand side of the editor, and click the arrow that points to the **Assigned Users** list.
3. To unassign, select the user or group in the **Assigned Users** list and click the arrow that points to the list of all users.

Caution

The double-arrow icon  unassigns all users and groups whether they are selected or not.

4. To save the changes in the repository, click the save icon in the main tool bar.

Related Information

[Opening the Security Editor](#) [page 335]

17.13 Displaying profiles assigned to a user and previewing net profiles

Procedure

1. In the Security Editor, click the **Users / Groups** pane on the left-hand side of the editor.
2. In the **Users / Groups** pane, select the user or group.
3. In the **Universes / Profiles** pane in the upper right-hand side of the editor, select the universe.

➔ Tip

You can change the display to list only universes that have profiles assigned to the selected user or group by selecting the **Display only universes assigned to the selected user/group** option.

Once you have selected a user and a universe, the assigned profiles appear in the profiles list on the lower right-hand side of the editor.

4. To preview the net Data Security Profile or net Business Security Profile, click **Preview Net Profile** below the corresponding profile list.

The Data Security Profile or Business Security Profile editor opens in read-only mode. The settings on each tab represent the settings that will be used after aggregation of all the profiles assigned to the user is taken into account.

Related Information

[Security profile aggregation](#) [page 348]











[Opening the Security Editor](#) [page 335]







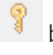

18 SQL and MDX reference

18.1 About the SQL/MDX Expression Editor

The SQL/MDX Expression Editor helps you write valid SQL/MDX expressions.

You can type SQL/MDX directly into the *Expression* box, or drag and drop table names, column names, business objects, functions, and parameters from the resource panes available in the editor. These panes are described in the following table. To display a resource pane, click the icon in the tool bar of the *Expression* pane. Different icons are available depending on the type of expression you are editing.

Icon	Description
 Tables	The list of tables and columns in the data foundation. To see a list of values for a column, click the  icon next to column name.
 Database Tables	For relational connections, the list of database tables in the connections. Used for defining expressions for derived tables and lists of values. To see a list of values for a column, click the  icon next to column name.
 OLAP Metadata	<p>For OLAP connections, the list of objects in the source cube.</p> <p>To change the display options, click . You can display names, keys, or both.</p> <p>To see a list of members for a level, click the  icon next to the level name. Use the  icon to search for a string in the object names.</p> <div><p> Note</p><p>A list of values for hierarchy attribute objects is not available in the OLAP Metadata pane. To see values for attributes, use the list in the Business Layer pane.</p></div>
 Functions	<p>The list of functions that can be used in the expression. The functions are grouped by type:</p> <ul style="list-style-type: none">• Operators: Common database operators, for example *, SUM, IS NOT NULL.• Database Functions: The SQL functions that are valid for the databases in the connections. For multisource-enabled data foundations or business layers, see the related topic about SAP BusinessObjects SQL functions.• System Variables: The system variables for which you can retrieve the values assigned using the @Variable function. For more information, see the related topic about @Variable.

Icon	Description
	<p> Note</p> <p>You can also reference User Attributes defined in the Central Management Server using @Variable.</p> <ul style="list-style-type: none"> • @Functions: The @Functions that are valid for this expression. For more information, see the related topic about @Functions.
 Business Layer	<p>The list of objects in the business layer. To see a list of members for a level, click the  icon next to the level name. Use the  icon to search for a string in the object names.</p> <p>Use the toggle button  to change how the object-related text is inserted into the expression:</p> <ul style="list-style-type: none"> • When the  button is unselected (default): The @Select function for the object is inserted, for example: @Select (Account\Account Number) • When the  button is selected: The SQL or MDX expression for the object is inserted, for example: [Account] . [Account Number]
 Parameters	<p>The list of parameters defined in the data foundation and business layer.</p>

Click the **Validate** icon in the tool bar of the [Expression](#) pane to check if the expression you have defined is valid SQL/MDX.

Date formats in SQL expressions

When entering a date value into an SQL expression, you need to use the format that is defined for each data source in the extended PRM file by the SQL generation parameter USER_INPUT_DATE_FORMAT.

For example, for ERP data sources, in the corresponding extended PRM file `jco.prm`, the parameter `USER_INPUT_DATE_FORMAT=DATE'yyyy-mm-dd'`. So, the SQL expression might look like this:

```
WHERE "table_name"."start_date"=DATE'2013-04-10'
```

For more information about SQL generation parameters and extended PRM files, see the related link.

Related Information

[SAP BusinessObjects SQL function reference for multisource-enabled universes](#) [page 360]

[About @Variable](#) [page 430]

[About @Functions](#) [page 423]

[About tables in the data foundation](#) [page 148]

[SQL generation parameters set in the extended PRM](#) [page 445]

18.2 SAP BusinessObjects SQL function reference for multisource-enabled universes

The information design tool provides a set of database functions based on SQL-92. Use these functions when defining SQL expressions for objects in a multisource-enabled data foundation or business layer.

This reference describes the syntax to use. The data federation service translates the SQL to the appropriate syntax for the data source at query run time.

i Note

The SAP BusinessObjects syntax can be different from the syntax of the same function provided by database-specific SQL.

18.2.1 Aggregation functions

18.2.1.1 Average (avg)

Description

Returns the average of a set of values.

Syntax

decimal avg(<set of values>)

Input

Parameter	Description	Data Type
<set of values>	A set of values.	Numeric

Notes

You can use the SQL keyword DISTINCT in front of column names.

Examples

Calculates the average of the sums of two columns: `avg(table.column1 + table.column2)`

Calculates the average of the values in a column that contains numbers written as strings:
`avg((toInteger(table.column1))`

18.2.1.2 Count

Description

Counts the number of values in a set.

Syntax

integer `count(<set of values>)`

Input

Parameter	Description	Data Type
<code><set of values ></code>	A set of values.	All data types (Numeric, String, Boolean, Date-Time, Date).

Notes

You can use the SQL keyword DISTINCT in front of column names.

Examples

Counts the number of values in a column: `count(table.column1)`

18.2.1.3 Maximum (max)

Description

Returns the maximum value in a set.

Syntax

value `max(<set of values>)`

Input

Parameter	Description	Data Type
<code><set of values></code>	A set of values.	All data types (Numeric, String, DateTime, Date).

Notes

You can use the SQL keyword DISTINCT in front of column names.

Examples

Returns the maximum value of a column: `max(table.column1)`

18.2.1.4 Minimum (min)

Description

Returns the minimum value in a set.

Syntax

value min(<set of values>)

Input

Parameter	Description	Data Type
set of values	A set of values.	All data types (Numeric, String, DateTime, Date).

Notes

You can use the SQL keyword DISTINCT in front of column names.

Examples

Returns the minimum value of a column: min(table.column1)

18.2.1.5 Sum

Description

Returns the sum of a set of values.

Syntax

decimal `sum(<set of values>)`

Input

Parameter	Description	Data Type
<code><set of values></code>	A set of values.	Numeric

Notes

You can use the SQL keyword DISTINCT in front of column names.

Examples

Sums the values in a column: `sum(table.column1)`

18.2.2 ASCII Code (`ascii`)

Description

Returns an integer representing the ASCII code value of the leftmost character in the input string.

Syntax

integer `ascii(<string>)`

Input

Parameter	Description	Data Type
<code><string></code>	A string of characters.	String

Notes

Returns null if **<string>** is null.

18.2.3 Absolute (abs)

Description

Returns the absolute value of a given integer value.

Syntax

numeric abs(**<expression>**)

Input

Parameter	Description	Data Type
<expression>	A numeric expression.	Numeric

Notes

- Returns null if the input **<expression>** is null.
- If **<expression>** is equal to the most negative value possible for an integer (-2 to the power of 31), that same negative value is returned.

18.2.4 Angle Tangent 2 (atan2)

Description

Returns the angle in radians whose tangent is **<angle1>/<angle2>**.

Syntax

numeric `atan2(<angle1>, <angle2>)`

Input

Parameter	Description	Data Type
<code><angle1 ></code>	An angle.	Numeric
<code><angle2></code>	An angle.	Numeric

Notes

Returns null if both `<angle1>` and `<angle2>` = 0.

Examples

`atan2(x,y)` converts rectangular coordinates (x, y) to polar (r, theta). This method computes the phase theta by computing an arc tangent of y/x in the range of -Pi to Pi.

18.2.5 Arc Tangent (atan)

Description

Returns the arc tangent of a given numeric expression.

Syntax

numeric `atan(<expression>)`

Input

Parameter	Description	Data Type
<code><expression></code>	An expression in the range from -Pi/2 to Pi/2.	Numeric

18.2.6 Arc Cosine (acos)

Description

Returns the arc cosine of a given numeric expression.

Syntax

numeric acos(`<expression>`)

Input

Parameter	Description	Data Type
<code><expression></code>	An expression in the range from 0 to Pi.	Numeric

Notes

Returns null if $\text{abs}(\text{acos}(\text{expression})) > 1$.

18.2.7 Arc Sine (asin)

Description

Returns the arc sine of a given numeric expression.

Syntax

numeric `asin(<expression>)`

Input

Parameter	Description	Data Type
<code><expression ></code>	An expression in the range from -Pi/2 to Pi/2.	Numeric

Notes

Returns null if `abs(<expression>) > 1`.

18.2.8 Case

Description

Returns a value depending on which of the given conditions is met.

Syntax

```
value CASE <input expression> WHEN <when expression> THEN <then result expression> ELSE  
<else result expression> END
```

Input

Parameter	Description	Data Type
<code><input expression></code>	An expression that represents a value to be compared to <code><when expression></code> .	All types. i Note <code><Input expression></code> must have the same

Parameter	Description	Data Type
		data type as <when expression> .
<when expression>	An expression that represents a value to be compared to <input expression> .	All types. i Note <Input expression> must have the same data type as <when expression> .
<then result expression>	An expression representing the value to return when <input expression> compared to <when expression> is true.	All types. i Note <Then result expression> must have the same data type as <else result expression> .
<else result expression>	An expression representing the value to return when <input expression> compared to <when expression> is not true.	All types. i Note <Then result expression> must have the same data type as <else result expression> .

Notes

- The `case` function implements the standard simple SQL `CASE` statement.
- `WHEN <when expression> THEN <then result expression>` is repeated to provide multiple conditions.

Examples

- `CASE (table1.column1)`

```

WHEN 'p1' THEN 'Product1'
WHEN 'p2' THEN 'Product2'
WHEN 'p3' THEN 'Product3'
ELSE 'Out of stock'
END

```

- CASE ProductName
WHEN 'laptop' THEN 1
ELSE 0
END

18.2.9 Cast

Description

Converts a given value to a given data type.

Syntax

value cast(<expression>, AS <data type>)

Input

Parameter	Description	Data Type
<expression>	A numeric expression.	All data types (Numeric, String, Boolean, Date-Time, Date).
<data type>	The data type to which to convert the value of <expression>.	A keyword that can have the following values: <ul style="list-style-type: none"> • NULL • VARCHAR • DOUBLE • DECIMAL • DATE • TIME • TIMESTAMP

18.2.10 Catalog

Description

Returns the default catalog of the connection.

Syntax

```
string catalog()
```

18.2.11 Ceil (ceiling)

Description

Returns the value of a number rounded up to the nearest integer.

Syntax

```
numeric ceiling(<expression>)
```

Input

Parameter	Description	Data Type
<expression>	A numeric expression.	Numeric

Notes

The type of the value returned is not converted. Therefore, `ceiling(1.9) = 2.0`. If you want to convert the value to an integer, use the conversion function `toInteger`.

18.2.12 Character (char)

Description

Returns the character corresponding to the given ASCII code.

Syntax

```
string char(<code>)
```

Input

Parameter	Description	Data Type
<code>	An ASCII code from 0 to 255.	Integer

Notes

Returns null if <code> < 0 or > 255.

18.2.13 Charindex (pos) (locate)

Description

Returns the position of a search string in a given character string.

Syntax

```
integer pos(<search string>,<string>,<start position>)
```

```
integer locate(<search string>,<string>,<start position>)
```

Input

Parameter	Description	Data Type
<code><search string></code>	The string you want to find the position of in <code><string></code> .	String
<code><string></code>	The string you want to search.	String
<code><start position></code>	The position in <code><string></code> where you want to start searching. If <code><start position></code> is not specified, the default start is position 1.	Integer

Notes

Returns 0 if the search string is not found.

Returns 0 if `<start position>` is longer than the length of `<string>`.

If `<start position>` \leq 0, the search starts at position 1.

Examples

```
pos('cd','abcd') = 3
```

```
pos('abc','abcd') = 1
```

```
pos('cd','abcdcd') = 3
```

```
pos('cd','abcdcd',3) = 3
```

```
pos('cd','abcdcd',4) = 5
```

```
pos('ef','abcd') = 0
```

18.2.14 Concat

Description

Concatenates two strings.

Syntax

```
string concat(<string1>, <string2>)
```

Input

Parameter	Description	Data Type
<string1>	A string.	String
<string2>	A string.	String

Notes

Returns null if either <string1> or <string2> is null.

Examples

```
concat('AB', 'CD') = 'ABCD'
```

18.2.15 Contains Only Digits

Description

Returns true (1) if the given string contains only digits. Otherwise the function returns false (0).

Syntax

```
boolean containsOnlyDigits(<string>)
```

Input

Parameter	Description	Data Type
<code><string></code>	A string.	String

18.2.16 `Convert`

Description

Converts a given value to a given data type.

Syntax

value `convert`(`<expression>`, `<data type>`)

Input

Parameter	Description	Data Type
<code><expression></code>	A value or expression.	All data types (Numeric, String, Boolean, Date-Time, Date).
<code><data type></code>	The data type to which to convert the value.	A string that can have the following values: <ul style="list-style-type: none">• NULL• INTEGER• DOUBLE• DECIMAL• DATE• TIME• TIMESTAMP

18.2.17 Cosine (cos)

Description

Returns the cosine of an angle.

Syntax

numeric cos(<angle>)

Input

Parameter	Description	Data Type
<angle>	An angle in radians.	Numeric

18.2.18 Cotangent (cot)

Description

Returns the cotangent of an angle in radians.

Syntax

numeric cot(<angle>)

Input

Parameter	Description	Data Type
<angle>	An angle in radians.	Numeric

Notes

Returns null if $\sin(\text{<angle>}) = 0$.

18.2.19 `Current Date (curDate)`

Description

Returns the current date.

Syntax

```
date curDate()
```

18.2.20 `Current Time (curTime)`

Description

Returns the current time.

Syntax

```
time curTime()
```

18.2.21 `Database`

Description

Returns the name of the database.

Syntax

string database()

18.2.22 Day Name

Description

Returns a string containing the day of the week of a given date.

Syntax

string dayName(<date>)

Input

Parameter	Description	Data Type
<date>	A date.	Date or DateTime

Notes

Returns the day name in English in uppercase letters. Possible values are as follows:

- SUNDAY
- MONDAY
- TUESDAY
- WEDNESDAY
- THURSDAY
- FRIDAY
- SATURDAY

18.2.23 Day Of Month

Description

Returns an integer from 1 to 31 representing the day of the month of a given date.

Syntax

```
integer dayOfMonth(<date>)
```

Input

Parameter	Description	Data Type
<date>	A date.	Date or DateTime

18.2.24 Day Of Week

Description

Returns an integer between 1 and 7 representing the day of the week of a given date. The first day of the week is Sunday.

Syntax

```
integer dayOfWeek(<date>)
```

Input

Parameter	Description	Data Type
<date>	A date.	Date or DateTime

18.2.25 Day Of Year

Description

Returns an integer between 1 and 366 representing the day of the year of a given date.

Syntax

```
integer dayOfYear(<date>)
```

Input

Parameter	Description	Data Type
<date>	A date.	Date or DateTime

18.2.26 Decrement Days

Description

Decrements a given date by the given number of days.

Syntax

```
date decrementDays(<date>, <number of days>)
```

Input

Parameter	Description	Data Type
<date>	A date.	Date or DateTime
<number of days>	The number of days to decrement the date.	Integer

18.2.27 Degrees

Description

Converts an angle measured in radians to an approximately equivalent angle measured in degrees.

Syntax

numeric degrees(<angle>)

Input

Parameter	Description	Data Type
<angle>	An angle in radians.	Numeric

18.2.28 Exp

Description

Returns the value of the mathematical constant e raised to the given exponent.

Syntax

numeric exp(<exponent>)

Input

Parameter	Description	Data Type
<exponent>	The exponential power.	Numeric

Examples

`exp(10)` = e to the power of 10 = 22,026.4658.

18.2.29 `Floor`

Description

Returns the value of a number rounded down to the nearest integer.

Syntax

`numeric floor(<expression>)`

Input

Parameter	Description	Data Type
<code><expression></code>	A numeric expression.	Numeric

Notes

The type of the value returned is not converted. Therefore, `floor(1.9)` = 1.0. If you want to convert the value to an integer, use the conversion function `toInteger`.

18.2.30 `Hexa To Int`

Description

Converts the hexadecimal value given by a string to integer.

Syntax

integer hexaToInt(<string>)

Input

Parameter	Description	Data Type
<string>	A string containing a hexadecimal value.	String

Examples

```
hexaToInt('AF') = 175
```

18.2.31 Hour

Description

Returns an integer from 0 to 23 representing the hour of a given time.

Syntax

integer hour(<time>)

Input

Parameter	Description	Data Type
<time>	A time.	DateTime

18.2.32 If Else

Description

Returns a value based on a given condition:

- If **<condition>** is true, the function returns the value of **<expression1>**.
- If **<condition>** is false, the function returns the value of **<expression2>**.

Syntax

```
value ifElse(<condition>, <expression1>, <expression2>)
```

Input

Parameter	Description	Data Type
<condition>	A logical expression.	Boolean
<expression1>	The value to return if <condition> resolves to true.	All data types (Numeric, String, Boolean, Date-Time, Date).
<expression2>	The value to return if <condition> resolves to false.	All data types (Numeric, String, Boolean, Date-Time, Date).

18.2.33 If Null (nvl)

Description

Returns a value based on whether or not a value is null:

- If **<expression1>** is null, the function returns the value of **<expression2>**.
- If **<expression1>** is not null, the function returns the value of **<expression1>**.

Syntax

```
value nvl(<value1>, <value2>)
```


Input

Parameter	Description	Data Type
<code><expression1></code>	Returns the value of <code><expression1></code> if that values is not null.	All data types (Numeric, String, Boolean, Date-Time, Date).
<code><expression2></code>	The value to return if <code><expression1></code> is null.	All data types (Numeric, String, Boolean, Date-Time, Date).

18.2.34 Increment Days

Description

Increments a given date by the given number of days.

Syntax

```
date incrementDays(<date>, <number of days>)
```

Input

Parameter	Description	Data Type
<code><date></code>	A date.	Date or DateTime
<code><number of days></code>	The number of days to increment the date.	Integer

18.2.35 Int To Hexa

Description

Converts a given integer to hexadecimal. The hexadecimal value is returned in a string.

Syntax

```
string intToHexa(<value>)
```

Input

Parameter	Description	Data Type
<value>	An integer.	Integer

Notes

- To ensure that the input value is of data type integer, you can use the `toInteger` function:
`intToHexa(toInteger (<value>)).`
- If `<value> < 0`, the function returns 'FFFFFFFF'.

18.2.36 Is Like

Description

Checks a string for a matching pattern. Returns true (1) if the function finds a match for the given pattern in the given string.

Syntax

```
boolean isLike(<string1>, <pattern>, <escape character>)
```

Input

Parameter	Description	Data Type
<string1>	A string.	String
<pattern>	A string containing the pattern you are trying to match in <string1>.	String

Parameter	Description	Data Type
	<p>The pattern can contain wildcard characters:</p> <ul style="list-style-type: none"> • The underscore character (_) matches any single character. • The percent sign character (%) matches any string of characters. <p>To match an underscore or percent sign in <string1>, define an escape character in <escape character> and precede the underscore or percent sign in <pattern> with the escape character.</p>	
<escape character> (optional)	A character that allows you to match wildcard characters in <string1> .	String

Notes

- Returns null if either **<string1>** or **<pattern>** is null.
- If **<escape character>** is specified and is null, returns null.
- If **<escape character>** is specified, every occurrence of the escape character in **<pattern>** must be followed by either an underscore or percent sign.

Examples

```
isLike('ABCD', 'AB%') = true
```

```
isLike('ABCD', 'AB_D') = true
```

```
isLike('10000', '100%') = true
```

```
isLike('10000', '100\%', '\') = false
```

```
isLike('status: 100%', '100\%', '\') = true
```

18.2.37 LPad

Description

Pads a string on the left side with a second given string to a given length.

Syntax

string lpad(<string1>, <string2>, <length>)

Input

Parameter	Description	Data Type
<string1>	A string.	String
<string2>	A string to insert into <string1> on the left side.	String
<length>	The total length of the return string after padding.	Integer

Notes

- If <length> < the length of <string1>, returns left(<string1>, <length>).
- Returns null if either <string2> is null or <length> <= 0.

18.2.38 Left

Description

Returns the given number of characters from the left of the given string.

Syntax

string left(<string>, <number of characters>)

Input

Parameter	Description	Data Type
<string>	A string.	String
<number of characters>	The number of leftmost characters to return.	Integer

Notes

Returns null if either `<string>` is null or `<number of characters>` is ≤ 0 .

18.2.39 Left Remove (ltrim)

Description

Removes the first sequence of spaces and tabs from the left side of the given string.

Syntax

```
string ltrim(<string>)
```

Input

Parameter	Description	Data Type
<code><string></code>	A string.	String

Examples

```
ltrim(' ABCD') = 'ABCD'
```

```
ltrim(' AB CD ') = 'AB CD '
```

18.2.40 Length

Description

Returns the length of a given string. Spaces are counted.

Syntax

integer length(<string>)

Input

Parameter	Description	Data Type
<string>	A string.	String

18.2.41 Log

Description

Returns the natural logarithm of the given value.

Syntax

double log(<expression>)

Input

Parameter	Description	Data Type
<expression>	A numeric expression > 0.	Double

Notes

Returns null if <expression> is <= 0.

18.2.42 `log10`

Description

Returns the common logarithm (base 10) of the given value.

Syntax

```
double log10(<expression>)
```

Input

Parameter	Description	Data Type
<expression>	A numeric expression > 0.	Double

Notes

Returns null if <expression> is <= 0.

18.2.43 `lowercase` (`lcase`)

Description

Converts a string to lowercase.

Syntax

```
string lcase(<string>)
```

Input

Parameter	Description	Data Type
<code><string></code>	A string.	String

Examples

```
lcase('ABCD') = 'abcd'
```

```
lcase('Cd123') = 'cd123'
```

18.2.44 Minute

Description

Returns an integer from 0 to 59 representing the minutes of a given date and time.

Syntax

```
integer minute(<time>)
```

Input

Parameter	Description	Data Type
<code><time></code>	A date and time.	DateTime

18.2.45 Mod

Description

Returns the remainder of the division of two integers: value1 / value2.

Syntax

integer `mod(<value1>, <value2>)`

Input

Parameter	Description	Data Type
<code><value1></code>	Value of the numerator.	Numeric
<code><value2></code>	Value of the divisor not equal to 0.	Numeric

Notes

Returns null if `<value2> = 0`.

18.2.46 `Month Name`

Description

Returns a string containing the month name of a given date.

Syntax

string `monthName(<date>)`

Input

Parameter	Description	Data Type
<code><date></code>	A date.	Date or DateTime

Notes

Returns the month name in English in uppercase letters. Possible values are as follows:

-
- JANUARY
 - FEBRUARY
 - MARCH
 - APRIL
 - MAY
 - JUNE
 - JULY
 - AUGUST
 - SEPTEMBER
 - OCTOBER
 - NOVEMBER
 - DECEMBER

18.2.47 Now

Description

Returns the current date and time.

Syntax

`dateTime now()`

18.2.48 Number of the Month (month)

Description

Returns an integer from 1 to 12 representing the month of a given date.

Syntax

`integer month(<date>)`

Input

Parameter	Description	Data Type
<code><date></code>	A date.	Date or DateTime

18.2.49 `Number of the Week (week)`

Description

Returns an integer from 1 to 53 representing the week in the year of a given date.

Syntax

integer `week(<date>)`

Input

Parameter	Description	Data Type
<code><date></code>	A date.	Date or DateTime

Notes

The first day of the week is Sunday. The first week in the year must contain at least one day. If January 1 is a Saturday, the following rules apply:

- January 1 is week 1.
- January 2 to 8 is week 2.
- December 25 to 31 is week 53.

18.2.50 `Permute`

Description

Permutes a given string using two templates: the `<reference template>` and the `<new template>`.

First, each character (or block of characters) in the **<reference template>** is assigned to a character (or block of characters) in the given string (**<string1>**). The lengths of **<string1>** and **<reference template>** must be the same.

Next, **<new template>** is used to permute the characters that were assigned in the **<reference template>**.

For example, the character string '22/09/1999', which represents a date, can be converted to '1999-09-22' as follows.

The **<reference template>** is 'DD/MM/YYYY'. The letters are assigned according to their position and grouping. So, 'DD' is the first block of characters, and is assigned the value '22', the first two characters in **<string1>**. The slash character (/) is assigned to the third character in **<string1>**. The next block of characters 'MM' is assigned to '09', and so on.

The **<new template>** is 'YYYY-MM-DD'. The permutation is applied and the resulting string is '1999-09-22'.

Text can also be inserted into **<new template>**, provided that none of the characters are already used in the **<reference template>**. For example, if **<new template>** = 'MM/DD Year: YYYY', the resulting string is '09/22 Year: 1999'.

Syntax

```
string permute(<string1>,<reference template>,<new template>)
```

Input

Parameter	Description	Data Type
<string1>	A string.	String
<reference template>	A string representing the pattern of <string1> .	String
<new template>	A string providing the new pattern for the permutation of <string1> .	String

Notes

- To represent a block of characters in the templates, repeat the character in the pattern. For example, 'YYYY' in **<reference template>** matches four characters in **<string1>**.
- The length of **<string1>** must be equal to the length of **<reference template>**, otherwise the function returns an error.

Examples

Change the format of how a date is represented:

- `permute('02/09/2003', 'DD/MM/YYYY', 'YYYY-MM-DD') = '2003-09-02'`
- `permute('02-09/200', 'DD/MM/YYYY', 'YYYY-MM-DD') = '2003-09-02'`
- `permute('02/09_2003', 'DD/MM/YYYY', 'DL :MM/DD An :YYYY') = 'DL :09/02 An :2003'`

Extract a month and year from a string of characters representing one date:

- `permute('2003-09-02', 'DDYY-MM-YY', 'MM/YY') = '09/03'`

Compose a number from an internal code:

- `permute('03/03/21-0123', 'bbbYY/MM/DD-NNNN', 'YYMMDDNNNN') = '0303210123'`

Extract date information from internal code:

- `permute('2003NL987M08J21', 'YYYYXXXXXXMMXDD', 'YYYY-MM-DD') = '2003-08-21'`

18.2.51 Pi

Description

Returns the constant value of Pi.

Syntax

`numeric pi()`

18.2.52 Power

Description

Returns the value of a number raised to the power of the given exponent.

Syntax

`numeric power(<value>, <exponent>)`

Input

Parameter	Description	Data Type
<value>	The base value.	Numeric
<exponent>	The exponent.	Integer

Notes

Returns null if <value> = 0 and <exponent> > 0.

18.2.53 Quarter

Description

Returns an integer from 1 to 4 representing the quarter of a given date. The value 1 represents January 1 through March 31.

Syntax

integer quarter(<date>)

Input

Parameter	Description	Data Type
<date>	A date.	Date or DateTime

18.2.54 Radians

Description

Converts an angle measured in degrees to an approximately equivalent angle measured in radians.

Syntax

numeric radians(<angle>)

Input

Parameter	Description	Data Type
<angle>	An angle in degrees.	Numeric

18.2.55 Random (rand)

Description

Returns a random number between 0 and 1. You can optionally provide a seed integer to initialize the random number generator.

Syntax

numeric rand(<value>)

Input

Parameter	Description	Data Type
<value> (optional)	Seed value for the random number generator.	Integer

18.2.56 Replace

Description

Replaces in a given string the occurrences of the pattern with a replacement string.

Syntax

```
string replace(<string>, <pattern>, <replacement string>)
```

Input

Parameter	Description	Data Type
<string>	A string.	String
<pattern>	The string of characters to search for and replace in <string>.	String
<replacement string>	The string to replace the <pattern> in <string>	String

Notes

- If <pattern> is null, returns <string>.
- Does not return null if <replacement string> is null.

Example

```
replace('rar', 'a', 'ada') = 'radar'
```

18.2.57 Replace String Exp

Description

Replaces in a given string all occurrences of the pattern with a replacement string following the syntax of a Java regular expression. For more information, refer to the Pattern documentation for Java regular expressions at <http://java.sun.com/j2se/1.4.2/docs/api/java/util/regex/Pattern.html>.

Syntax

```
string replaceStringExp(<string>, <pattern>, <replacement string>)
```


Input

Parameter	Description	Data Type
<code><string></code>	A string.	String
<code><pattern></code>	The string of characters to search for and replace in <code><string></code> .	String
<code><replacement string></code>	The string to replace the <code><pattern></code> in <code><string></code>	String

18.2.58 `Replicate` (`repeat`)

Description

Returns a string formed by repeating a given character string a given number of times.

Syntax

```
string repeat(<string>,<number of replications>)
```

Input

Parameter	Description	Data Type
<code><string></code>	A string.	String
<code><number of replications></code>	The number of times to repeat <code><string></code> .	Integer

Notes

Returns null if `<number of replications>` ≤ 0 .

18.2.59 Rightpart (right)

Description

Returns the given number of characters from the right of the given string.

Syntax

```
string right(<string>, <number of characters>)
```

Input

Parameter	Description	Data Type
<string>	A string.	String
<number of characters>	The number of rightmost characters to return.	Integer

Notes

Returns null if either <string> is null or <number of characters> is <= 0.

18.2.60 Round

Description

Returns a number rounded to the given number of decimal places.

Syntax

```
numeric round(<expression>, <number of places>)
```

Input

Parameter	Description	Data Type
<code><expression></code>	The base value to be rounded.	Numeric
<code><number of places></code>	The number of decimal places to round to.	Numeric

Notes

- The function rounds towards the nearest neighboring integer unless both neighbors are equidistant, in which case it rounds away from zero.
- The type of the value returned is not converted. Therefore, `round(1.9) = 2.0`. If you want to convert the value to an integer, use the conversion function `toInteger`.

18.2.61 `Rpad`

Description

Pads a string on the right side with a second given string to a given length.

Syntax

```
string rpad(<string1>, <string2>, <length>)
```

Input

Parameter	Description	Data Type
<code><string1></code>	A string.	String
<code><string2></code>	A string to insert into <code><string1></code> .	String
<code><length></code>	The total length of the return string after padding.	Integer

Notes

- If **<length>** < the length of **<string1>**, returns `right(<string1>, <length>)`.
- Returns null if either **<string2>** is null or **<length>** <= 0.

18.2.62 rpos

Description

Returns the position of the last occurrence of a search string in a given character string.

Syntax

integer `rpos(<search string>, <string>)`

Input

Parameter	Description	Data Type
<search string>	The string you want to find the last occurrence of in <string> .	String
<string>	The string you want to search.	String

Notes

Returns 0 if the search string is not found.

Examples

`rpos('cd','abcd') = 3`

`rpos('cd', 'abcdcd') = 5`

`rpos('abc', 'abcdcd') = 1`

`rpos('ef', 'abcd') = 0`

18.2.63 rtrim

Description

Removes the first sequence of spaces and tabs from the right side of the given string.

Syntax

```
string rtrim(<string>)
```

Input

Parameter	Description	Data Type
<string>	A string.	String

Examples

```
rtrim('ABCD ') = 'ABCD'
```

```
rtrim(' AB CD ') = ' AB CD'
```

18.2.64 schema

Description

Returns the default schema (qualifier and owner) of the current connection.

Syntax

```
string schema()
```

18.2.65 Second

Description

Returns an integer from 0 to 59 representing the seconds of a given date and time.

Syntax

integer second(<time>)

Input

Parameter	Description	Data Type
<time>	A date and time.	DateTime

18.2.66 Sign

Description

Returns the positive (1), zero (0), or negative (-1) sign of a given number.

Syntax

numeric sign(<value>)

Input

Parameter	Description	Data Type
<value>	A numeric value.	Integer

18.2.67 Sine (sin)

Description

Returns the sine of an angle.

Syntax

numeric `sin(<angle>)`

Input

Parameter	Description	Data Type
<code><angle></code>	An angle in radians.	Numeric

18.2.68 Space

Description

Returns a character string with the given number of blank characters (spaces).

Syntax

string `space(<number of spaces>)`

Input

Parameter	Description	Data Type
<code><number of spaces></code>	The number of spaces to be returned in the string.	Integer

Notes

Returns null if **<number of spaces>** <=0.

18.2.69 Sqrt

Description

Returns the square root of a number.

Syntax

numeric sqrt(**<expression>**)

Input

Parameter	Description	Data Type
<expression>	An expression >= 0.	Numeric

Notes

Returns null if expression < 0.

18.2.70 Stuff (insert)

Description

Replaces the sequence of characters in one given string with a second given string.

Syntax

string insert(**<string1>**, **<start position>**, **<number of characters>**, **<string2>**)

Input

Parameter	Description	Data Type
<code><string1></code>	A string.	String
<code><start position></code>	A number representing the position in <code><string1></code> to start the replacement. Must be in the range of 1 to the length of <code><string1></code> + 1.	Integer
<code><number of characters></code>	The number of characters in <code><string1></code> to replace. Must be in the range of 0 to the length of <code><string1></code> .	Integer
<code><string2></code>	The replacement string.	String

Notes

Returns null if either `<start position>` or `<number of characters>` is out of range.

18.2.71 Substring

Description

Returns a substring of a given string.

Syntax

string substring(`<string>`, `<start position>`, `<number of characters>`)

Input

Parameter	Description	Data Type
<code><string></code>	A string.	String
<code><start position></code>	The start position in <code><string></code> of the substring.	Integer

Parameter	Description	Data Type
	Must be in the range of 1 to the length of <string> .	
<number of characters>	The number of characters to include in the substring.	Integer

Notes

Returns null in the following situations:

- **<start position>** <= 0
- **<start position>** > the length of **<string>**
- **<string>** is null
- **<number of characters>** <= 0

Examples

```
substring('ABCD', 2, 2) = 'BC'
```

```
substring('ABCD', 2, 10) = 'BCD'
```

```
substring('ABCD', 0, 2) = null
```

18.2.72 Tangent (tan)

Description

Returns the tangent of an angle.

Syntax

```
numeric tan(<angle>)
```

Input

Parameter	Description	Data Type
<code><angle></code>	An angle in radians.	Numeric

Notes

Returns null if $\cos(\text{<angle>}) = 0$.

18.2.73 Timestamp Add

Description

Returns a timestamp calculated by adding the given number of intervals to the given timestamp.

Syntax

`dateTime timestampAdd(<interval>, <count>, <timestamp>)`

Input

Parameter	Description	Data Type
<code><interval ></code>	An interval constant. This parameter can either be a string or integer constant as follows: <ul style="list-style-type: none">• 'SQL_TSI_FRAC_SECOND' or 0• 'SQL_TSI_SECOND' or 1• 'SQL_TSI_MINUTE' or 2• 'SQL_TSI_HOUR' or 3• 'SQL_TSI_DAY' or 4• 'SQL_TSI_WEEK' or 5• 'SQL_TSI_MONTH' or 6• 'SQL_TSI_QUARTER' or 7• 'SQL_TSI_YEAR' or 8	String or integer
<code><count ></code>	The number of intervals to add to the timestamp.	Integer

Parameter	Description	Data Type
<timestamp>	A date and time.	DateTime

Notes

The calculation can be impacted by daylight saving time in the locale for 'SQL_TSI_HOUR'.

18.2.74 Timestamp Diff

Description

Returns an integer representing the number of intervals by which the first given timestamp is greater than the second given timestamp.

Syntax

```
integer timestampDiff(<interval>, <timestamp1>, <timestamp2>)
```

Input

Parameter	Description	Data Type
<interval>	An interval constant. This parameter can either be a string or integer constant as follows: <ul style="list-style-type: none"> 'SQL_TSI_FRAC_SECOND' or 0 'SQL_TSI_SECOND' or 1 'SQL_TSI_MINUTE' or 2 'SQL_TSI_HOUR' or 3 'SQL_TSI_DAY' or 4 'SQL_TSI_WEEK' or 5 'SQL_TSI_MONTH' or 6 'SQL_TSI_QUARTER' or 7 'SQL_TSI_YEAR' or 8 	String or integer
<timestamp1>	A date and time.	DateTime
<timestamp2>	A date and time.	DateTime

Notes

- The calculation can be impacted by daylight saving time in the locale for 'SQL_TSI_HOUR'.
- Large differences can cause an error.
- The first day of the week is Sunday.

18.2.75 To Boolean

Description

Converts a given value to a Boolean value.

Syntax

boolean toBoolean(<expression>)

Input

Parameter	Description	Data Type
<expression>	A value or expression.	String or Boolean

Examples

```
toBoolean('true') = 1
```

```
toBoolean('TrUe') = 1
```

```
toBoolean('tru') = 0
```

```
toBoolean('False') = 0
```

```
toBoolean('F') = 0
```

```
toBoolean('f') = 0
```

18.2.76 To Date

Description

Converts a character string to a date.

Syntax

date toDate(<string>)

Input

Parameter	Description	Data Type
<string>	A string containing a date value in the format: yyyy-mm-dd, where yyyy is the year, mm is the month, and dd is the day. For example, 2003-09-07 and 2003-11-29.	String

Notes

- If <string> does not use the correct format, an error is returned.
- No restrictions are imposed on the month, day, or year values. If the month is greater than 12 or the day does not exist in the corresponding month, the function uses the internal calendar to convert to the correct date.

Examples

toDate('2003-02-12') = February 12, 2003

toDate('2003-02-29') = March 1, 2003

toDate('2002-14-12') = February 12, 2003

toDate('1994-110-12') = February 12, 2003

18.2.77 To Decimal

Description

Converts a given value to a decimal.

Syntax

decimal toDecimal(<expression>)

Input

Parameter	Description	Data Type
<expression>	A value. If the value is a string, it must be in decimal number format and use the period character (.) as the decimal separator.	Numeric or String

18.2.78 To Double

Description

Converts a given value to a decimal.

Syntax

double toDouble(<expression>)

Input

Parameter	Description	Data Type
<expression>	A value.	Numeric or String

Parameter	Description	Data Type
	If the value is a string, the input must be in decimal number format and use the period character (.) as the decimal separator.	

18.2.79 To Integer

Description

Converts a given value to an integer.

Syntax

integer toInteger(<expression>)

Input

Parameter	Description	Data Type
<expression>	A value. If the value is a string, the input must be in number format.	Numeric or String

18.2.80 To Null

Description

Converts a given value to null.

Syntax

null toNull(<expression>)

Input

Parameter	Description	Data Type
<code><expression></code>	A value.	All data types (Numeric, String, Boolean, Date-Time, Date).

18.2.81 To String

Description

Converts a given value to a string.

Syntax

`string toString(<expression>)`

Input

Parameter	Description	Data Type
<code><expression></code>	A value.	All data types (Numeric, String, Boolean, Date-Time, Date).

Examples

`toString(45) = '45'`

`toString(-45) = '-45'`

`toString(45.9) = '45.9'`

`toString(-45.9) = '-45.9'`

`toString(Date value for September 9, 2002) = '2002-09-09'`

`toString(DateTime value for September 9, 2002 23:08:08) = '2002-03-03 23:08:08'`

`toString(Boolean value 1) = 'true'`

toString(Boolean value 0) = 'false'

18.2.82 To Time

Description

Converts a given value to a time.

Syntax

time toTime(<expression>)

Input

Parameter	Description	Data Type
<expression>	A value. If the value is a string, the input must be in the format: hh:mm:ss where hh is the hour, mm is the minutes, and ss is the seconds. For example, 23:09:07 and 03:11:23.	String, Date, Time, or DateTime

Notes

- If <expression> does not use the correct format, an error is returned.
- No restrictions are imposed on the hour, minutes, or seconds values. If the minutes or seconds are greater than 60, or if the hour is greater than 24, the function uses the internal clock to convert to the correct time.

Examples

toTime('02:10:09') = '02:10:09'

toTime('0:450:29') = '07:30:29'

toTime('25:14:180') = '01:17:00'

18.2.83 To Timestamp

Description

Converts a given value to a date and time.

Syntax

time toTimestamp(<expression>)

Input

Parameter	Description	Data Type
<expression>	<p>A value.</p> <p>If the value is a string, the input must be in the format: yyyy-mm-dd hh:mm:ss.ssss, where yyyy is the year, mm is the month, dd is the day, hh is the hour, mm is the minutes, ss is the seconds, and ssss is the milliseconds (optional).</p> <p>For example, 2003-09-07 23:09:07 and 2003-11-29 03:11:23.0.</p>	String, Date, Time, or DateTime

Notes

- If <expression> does not use the correct format, an error is returned.
- No restrictions are imposed on the month, day, or year values. If the month is greater than 12 or the day does not exist in the corresponding month, the function uses the internal calendar to convert to the correct date.
- No restrictions are imposed on the hour, minutes, or seconds values. If the minutes or seconds are greater than 60, or if the hour is greater than 24, the function uses the internal clock to convert to the correct time.

Examples

```
toTimestamp('2003-02-12 02:10:09') = '2003-02-12 02:10:09.0'
```

```
toTimestamp('2003-02-29 02:10:09') = '2003-03-01 02:10:09.0'
```

```
toTimestamp('2002-14-12 02:10:09') = '2003-02-12 02:10:09.0'
```

```
toTimestamp('1994-11-12 02:10:09') = '2003-02-12 02:10:09.0'
toTimestamp('2003-02-12 0:45:29') = '2003-02-12 07:30:29.0'
toTimestamp('2002-09-09 25:14:180') = '2002-09-09 01:17:00.0'
```

18.2.84 Trim

Description

Removes the spaces and tabs from the left and right sides of the given string.

Syntax

```
string trim(<string>)
```

Input

Parameter	Description	Data Type
<string>	A string.	String

18.2.85 Trunc

Description

Returns a number truncated to a given number of decimal places.

Syntax

```
numeric trunc(<expression>, <number of places>)
```

Input

Parameter	Description	Data Type
<expression>	The base value to be truncated.	Decimal
<number of places>	The number of decimal places to remain after truncating.	Integer

Notes

- If <number of places> is omitted, the number is truncated to 0 decimal places.
- If <number of places> is negative, the function starts at the digit the number of places to the left of the decimal point and sets to zero all the digits to the right of that position.

Examples

`trunc(10.1234, 1) = 10.1`

`trunc(10.1234, 2) = 10.2`

`trunc(1862.1234, -1) = 1860`

`trunc(1862.1234, -2) = 1800`

18.2.86 Uppercase (ucase)

Description

Converts a string to uppercase.

Syntax

`string ucase(<string>)`

Input

Parameter	Description	Data Type
<code><string></code>	A string.	String

Examples

```
ucase('abcd') = 'ABCD'
```

18.2.87 `User`

Description

Returns the user name as defined in the connection parameters.

Syntax

```
string user()
```

18.2.88 `Year`

Description

Returns an integer representing the year of a given date.

Syntax

```
integer year(<date>)
```

Input

Parameter	Description	Data Type
<code><date></code>	A date.	Date or DateTime

18.3 About @Functions

@Functions are special functions that provide more flexible methods for specifying the query script for an object. Select the related topic to see more about an @Function.

Related Information

[About @Aggregate_Aware](#) [page 423]

[About @DerivedTable](#) [page 424]

[About @Execute](#) [page 424]

[About @Prompt](#) [page 426]

[About @Select](#) [page 430]

[About @Variable](#) [page 430]

[About @Where](#) [page 432]

18.3.1 About @Aggregate_Aware

Use the @Aggregate_Aware function in the SQL definition of a business layer object to make the object aggregate-aware. When the object is included in a query, the aggregate tables listed as parameters in the @Aggregate_Aware function are queried first.

The syntax is:

```
@Aggregate_Aware(sum(<Aggregate table 1>), ... sum(<Aggregate table n>))
```

<Aggregate table 1> is the aggregate table with the highest level of aggregation, and **<Aggregate table n>** is the aggregate table with the lowest level.

For more information about aggregate awareness in the universe, see the related topic.

Related Information

[About aggregate awareness](#) [page 241]

18.3.2 About @DerivedTable

Use the @DerivedTable function in the definition of nested derived tables. A nested derived table (also known as a 'derived table on a derived table') is a table that is derived from at least one existing derived table.

The syntax of the @DerivedTable function is:

```
@DerivedTable(<Derived table name>)
```

<Derived table name> is the name of the derived table you want to reference. The @DerivedTable function is only used in the definition of derived tables in the data foundation.

Note

In database-specific SQL (multisource-enabled data foundations), all tables referenced must be from the same connection.

18.3.3 About @Execute

The @Execute function lets you define a preliminary query that provides a list of values in a SELECT predicate to be included in the main query. The @Execute function is based on standard SQL and so applies to relational data sources. The syntax of the @Execute function is:

```
@Execute(<List of values>)
```

<List of values> is a list of values pre-defined in the business layer or data foundation. The list of values definition provides the preliminary query. Most often the @Execute function is then included in a filter or WHERE clause to apply the preliminary query to limit the values returned in the main query.

The list of values can be any of the following types:

- List of values based on custom SQL
- Static list of values
- List of values based on a query that includes business layer objects

The following limitations apply:

- The list of values cannot be based on a custom hierarchy.
- The list of values can contain only objects that are active in the business layer (not hidden or deprecated).
- The SQL defining the list of values cannot contain the @Execute function.
- The @Execute function cannot be used in the definition of an @Prompt function.

For more information on inserting a list of values, see the related topic.

Example

Filter on products

This example creates a query filter that limits query results to products with sales that are two times above the product category average.

First, create the list of values that returns the product IDs of products with sales above the category average. The list of values name is **Products_Above_Avg**, and the data type is numeric. The following SQL defines the list of values:

```
WITH
PA as
(
  SELECT L.PRODUCT_ID, sum(L.NET_SALES) AS SALES
  FROM PRODUCT P, PA A
  FROM SO_LINE L
  GROUP BY L.PRODUCT_ID
),
CA as
(
  SELECT P.CATEGORY_ID, avg(A.SALES) AS
  SALES
  WHERE P.PRODUCT_ID = A.PRODUCT_ID
  GROUP BY P.CATEGORY_ID
)
SELECT PA.PRODUCT_ID
FROM PA, CA, PRODUCT P
WHERE PA.PRODUCT_ID = P.PRODUCT_ID
AND P.CATEGORY_ID = CA.CATEGORY_ID
AND PA.SALES > ( CA.SALES * 2)
```

Next, insert a native filter into the business layer that invokes the preliminary query using the @Execute function in the WHERE clause. Because the @Execute function can return multiple values, use the IN operator in the filter definition:

```
PRODUCT.PRODUCT_ID IN
@Execute(Products_Above_Avg)
```

When the filter is included in a query, the @Execute function is replaced with the resulting list of product IDs, for example:

```
PRODUCT.PRODUCT_ID in (2, 5, 20, 33, 35)
```

Example

Include a security predicate

This example inserts a column filter that returns sales data only for the geographical region of the current user.

First, create the list of values in the data foundation that returns the authorized country codes for the current user. The list of values name is **Authorized_Countries**, and the data type is numeric. This example assumes the database administrator has set up a table called **user_geography** in the database that associates authorized countries with each user. The following SQL defines the list of values:

```
SELECT country_id
FROM user_geography
WHERE user_name = @Variable('BOUSER')
```

Next, insert a column filter into the data foundation table **Sales**. Because the @Execute function can return multiple values, use the IN operator in the filter definition.

```
Sales.country_id
```

```
IN @Execute(Authorized_Countries)
```

When a user includes the **Sales** table in a query, the @Execute function in the column filter is replaced by the list of authorized country codes for that user.

Related Information

[Inserting or editing a list of values](#) [page 284]

[Inserting and editing filters](#) [page 255]

[Inserting a column filter](#) [page 164]

18.3.4 About @Prompt

Use the @Prompt function to insert a prompt into a query. Prompts can be used to restrict the data when a user creates a report. You use the @Prompt function in the SQL SELECT statement or WHERE clause, or in the MDX expression for an object. It forces a user to enter one or more values (or select from a list of values) for a restriction when that object is used in a query. When the user runs the query, a prompt box appears asking for a value to be entered or selected.

Prompts are useful when you want to force a restriction in the query script, but do not want to preset the value of the condition.

The @Prompt function is allowed in the following expressions:

- Joins
- Calculated columns (except in database-specific SQL in multisource-enabled data foundations)
- Derived tables
- Business objects in the business layer

You can insert a @Prompt definition in the following ways:

- Define a named parameter for the prompt and reference the parameter in the @Prompt function, for example:
@Prompt(**<Parameter name>**)
<Parameter name> is a parameter pre-defined in the data foundation or business layer. For more information, see the related topic about parameters.
- Type the prompt definition into the SQL or MDX expression of the object. For more information on the syntax and parameters of the @Prompt function, see the related topic.

Related Information

[About parameters](#) [page 280]

[@Prompt syntax](#) [page 427]

18.3.4.1 @Prompt syntax

The syntax for the @Prompt function is as follows:

```
@Prompt('<message>',  
'<type>',  
'<folder\business layer object>' | '<list of values>' | {'<value_1>','<value_2>',...},  
Mono | Multi : Any | Leaf,  
free | constrained | primary_key,  
persistent | not_persistent,  
{'<default_value_1>',... '<default_value_n>'})
```

The function parameters and the possible values are described in the following table. The parameters are separated by commas. You must specify at least the first two parameters. If you want to specify additional parameters, you must include the intervening commas for the optional parameters.

Parameter	Description
'<message>'	<p>Text of the prompt message. This parameter is mandatory.</p> <p>The text appears in the prompt box when the user runs the query.</p> <p>The text must be enclosed in single quotes, for example, 'Choose a Region'.</p> <p>For proper function of the prompt, the prompt text should be unique within the universe.</p>
'<type>'	<p>The data type of the prompt. This parameter is mandatory.</p> <p>The user's response is interpreted with the data type you specify. The list of values and default values also have this data type. It can be one of the following:</p> <ul style="list-style-type: none">• 'A' for alphanumeric string.• 'K' for keyword. This type is also an alphanumeric string, however the responses to the prompt will be not be surrounded by quotes in the query script at run time.• 'N' for number.• 'D' for date.• 'DT' for date-time. <p>The specified <type> must be enclosed in single quotes.</p> <p>The <type> parameter can be a pair of data types to indicate a name and key. The syntax is '<name_type>:<key_type>', for example: 'A': 'N' where the first type is the data type of the name that the user sees in the list of values, and the second type is the data type of the primary key that is used by the query. Both the <name_type> and the <key_type> can be any of the available data types.</p>

Parameter	Description
	<p>i Note</p> <p>To use this option, you must make sure the object and the prompt are index aware:</p> <ul style="list-style-type: none"> Define a primary key for the object in the business layer. Specify primary key for the fifth parameter in the @Prompt function. <p>In this case, if the list of values or default values parameters are used, they must contain a list of pairs of values.</p>
<p>'<folder\business layer object>' </p> <p>'<list of values>' </p> <p>{'<value_1>','<value_2>',...}</p>	<p>The list of values the user can choose from when prompted. This parameter is optional.</p> <p>You can specify a list of values in three ways:</p> <ul style="list-style-type: none"> The default list of values associated with an object in the business layer (a dimension, measure, attribute, hierarchy, or hierarchy level). Enter the full path and object name in the business layer between single quotes, for example: 'Myconnection\dimproduct\productname' In this example, productname is the business layer object name. The object must be index aware, that is, a primary key is defined for the object in the business layer. For more information, see the related topic on defining keys. A named list of values defined in the business layer or data foundation. Enter the name of the list of values between single quotes, for example: 'G7_Countries'. If the list of values is hierarchical with named levels, you can specify the level to use for the prompt, for example: 'Country_Region_City_List':'Region' In this example, Country_Region_City_List is the name of the list of values and Region is the target level. If the list of values is a multiple column list of values with named columns, you can specify the column to use for the prompt, for example: 'Country_Region_City_List':'Region'. In this example, Country_Region_City_List is the name of the list of values and Region is the target column. A hard-coded list of values or name/key pairs. The values in a pair are separated by a colon. Each value is enclosed in single quotes. The pairs of values are separated by a comma. The whole list is enclosed in braces: The syntax for a single value: {'<value>'} The syntax for several single values: {'<value_1>','<value_2>',...,'<value_n>'} <p>The syntax for a pair of values: {'<name_value>':'<key_value>'}</p> <p>The syntax for pairs of values: {'<name_value_1>':'<key_value_1>','<name_value_2>':'<key_value_2>'}</p>

Parameter	Description
	<p>'...,<name_value_n>':<key_value_n>'. For example: {'France':'FR', 'Germany':'DE', 'Spain':'ES', 'United Kingdom':'UK'}</p> <div> <p>i Note</p> <p>If the list of values is index aware (a primary key has been defined for the object in the business layer, or you use {name, key} pairs for the list of values, specify primary key for the fifth parameter in the @prompt function.</p> </div>
Mono Multi : Any Leaf	<p>The selection mode. If not specified, Mono is default.</p> <ul style="list-style-type: none"> Use Mono if the user can select only one value from the list of values. Use Multi if the user can select multiple values from the list of values. <p>You can optionally specify the hierarchical selection mode for hierarchical lists of values. If not specified, Leaf is default:</p> <ul style="list-style-type: none"> Use Any if the user can select any member/value at any level of the hierarchical list of values. Use Leaf if the user can select only the leaf members/values from the hierarchical list of values.
free constrained primary_key	<p>The entry constraint type. If not specified, free is default.</p> <ul style="list-style-type: none"> Use free if the user can enter a value, or select from the list of values. Use constrained if the user must select values from the list of values. Use primary_key when you are using an index aware object or {name, key} pairs. The associated key value for the object is used in the query rather than the entered or displayed name value.
persistent not_persistent	<p>Whether or not the last values are displayed. If not specified, not_persistent is default.</p> <p>Use persistent if, when refreshing a document, the last values used in the prompt are displayed by default, even when default values are defined.</p> <p>Use not_persistent if, when refreshing a document, no values used are displayed in the prompt by default.</p>
{'<default value>' }	<p>One or more default values presented to the user. This parameter is optional.</p> <p>Enter default values in the following ways:</p> <ul style="list-style-type: none"> For a single value: {'France'} For a pair of values: {'France':'FR'} For two pairs of values: {'France':'FR','Germany':'DE'} For hierarchical values, use \ to separate the hierarchy level values: {'Europe':'2\'France\'Marseille\'CSP Systems','Europe':'2\'Germany\'Berlin'}

Parameter	Description
	<p>When refreshing a document, these values are displayed by default, but if the persistent option is set, then the last values used in the prompt are used instead of the default values.</p> <p>If you specify the primary_key parameter in the prompt definition, you must provide the key values.</p>

Related Information

[Defining keys for dimensions and dimension attributes](#) [page 247]

18.3.5 About @Select

Use the @Select function in the definition of an object in the business layer to re-use the SELECT statement of another object. The syntax of the @Select function is:

```
@Select(<Folder name>\<Object name>)
```

<Folder name>\<Object name> specifies the full path of another object in the business layer.

For example, you define a business layer object **Promotional_Service_Line** as @Select(**Resort \Service_Line**). The SELECT statement defined for **Service_Line** is used for the definition of **Promotional_Service_Line**.

Using the @Select function lets you maintain only one instance of the SQL or MDX expression, and ensures the consistency of related object definitions in the business layer. However, @Select creates an object dependency. If you delete the source object, you must manually update the object that uses the @Select function.

18.3.6 About @Variable

Use the @Variable function in an SQL or MDX expression (usually in the WHERE clause) to retrieve the value assigned to a system variable or User Attribute. The syntax of the @Variable function is:

```
@Variable('<Variable name>' [, DELIMITER=default | no_quote])
```

<Variable name> must be in single quotes. Possible variables are described in the following table:

Variable Name and Description	Examples
<p>Variables containing information about the user's authorization:</p> <ul style="list-style-type: none"> BOUSER: User name entered by the user to log into the SAP BusinessObjects BI platform. DBUSER: User name used for authorization when connecting to the data source. This user name can be defined in the Central Management Console as part of the user's secondary credentials. 	<p>For example, to restrict data retrieved in a query to the current user, use the BOUSER variable in the WHERE clause:</p> <pre>WHERE Employees.Employee_Name = @Variable('BOUSER')</pre>
<p>Variables containing information about the current report or query:</p> <ul style="list-style-type: none"> DOCNAME: The document name. DOCID: The document identifier. (If the document is published in the repository, the DOCID value corresponds to the document ID in the repository. If the document is not published in the repository, the DOCID value is EMPTY.) DPNAME: The Data Provider name. DPTYPE: The Data Provider type. UNVNAME: The universe name. UNVID: The universe identifier. 	<p>For example, these variables can be referenced in the BEGIN_SQL parameter that will be executed before the SELECT statement. This can be used for audit purposes concerning the use of the database (For example: To determine which report query or which universe is used most frequently).</p>
<p>Variables containing information about the user's current language settings:</p> <ul style="list-style-type: none"> PREFERRED_VIEWING_LOCALE: The user's preferred locale for viewing report and query objects in an application. DOMINANT_PREFERRED_VIEWING_LOCALE: A pre-defined fallback locale that is used when no fallback locale is defined for the resource. 	<p>The following query retrieves the product names in the language determined by the user's Preferred Viewing Locale. The database must contain a column identifying the locale of the data. To see a list of locales, their abbreviations and dominant locales, refer to the <i>Translation Management Tool User Guide</i>.</p> <pre>SELECT Product_Name FROM Product WHERE Product.Locale = @Variable('PREFERRED_VIEWING_LOCALE')</pre>
<p>User Attributes defined in the User Attribute Management area of the Central Management Console (CMC).</p>	<p>To reference a User Attribute, specify the internal name for the attribute as it is defined in the CMC. @Variable returns the value of the attribute for the current user. For example, the User Attribute MYCOUNTRY contains the value of the country of each user in the CMC. Specify the attribute's internal name surrounded by single quotes:</p> <pre>@Variable('SI_MYCOUNTRY')</pre>

Variable Name and Description	Examples
	The attribute's internal name is defined when the attribute is created in the CMC.

i Note

If the **<variable name>** specified in the @Variable function is unknown to the system, the user is prompted for a value. In this case, the @Variable function behaves in the same way as a single-value @Prompt function with the following settings:

```
@Prompt('<Variable name>','A',,Mono,free)
```

The DELIMITER parameter specifies how the returned value for the variable is delimited in the query script. The parameter's default value is DELIMITER=default. This means the value is delimited by quotes for relational SQL data sources, and no delimiter for OLAP MDX data sources.

If you specify DELIMITER=no_quote, this means that no delimiter character is added around the value in the script.

The @Variable function is allowed in the following expressions:

- Joins
- Calculated columns
- Derived tables
- Object definitions in the business layer
- BEGIN_SQL and END_SQL statements
- Connection properties, such as the ConnectInit property (except the DELIMITER parameter, which is not supported in connection properties)

In database-specific SQL (multisource-enabled data foundations), all referenced tables or columns must be from the same connection.

18.3.7 About @Where

Use the @Where function in the SQL definition of an object in the business layer to re-use the WHERE clause of another object. The syntax of the @Where function is:

```
@Where(<Folder name>\<Object name>)
```

<Folder name>\<Object name> specifies the full path of another object in the business layer.

For example, you define the WHERE clause of business layer object **Resort_Service_Line** as @Where(**dimResort\Resort**). The WHERE statement defined for the object **Resort** is used for the definition of **Resort_Service_Line**.

Using the @Where function lets you maintain only one instance of the SQL WHERE clause, and ensures the consistency of related object definitions in the business layer. However, @Where creates an object dependency. If you delete the source object, you must manually update the object that uses the @Where function.

18.4 About SQL Generation Parameters

SQL generation parameters affect the generation of the query script. The parameters all have default values. Default values can be overridden in the data foundation properties. Some parameters (concerning lists of values) can also be overridden in the business layer properties. At query time, the query server will use the values it finds in the following order:

1. The value in the business layer, if it is set.
2. The value in the data foundation, if it is set.
3. The default value.

The following reference describes the parameters that affect the generation of the query script. The parameters are listed alphabetical order in two groups:

- SQL parameters that you set in the user interface of the information design tool. These are SQL parameters that are common to most data access drivers. Each parameter is valid for the universe in which it is set.
- SQL parameters that you set in the extended data access parameter (PRM) files. These are connection-specific parameters that are listed in the extended PRM file for the target data access driver.

Related Information

[SQL generation parameters reference](#) [page 433]

[SQL generation parameters set in the extended PRM](#) [page 445]

[About data foundation properties](#) [page 180]

[About business layer properties](#) [page 227]

18.4.1 SQL generation parameters reference

The following reference describes the SQL generation parameters that can be overridden in the data foundation properties and business layer properties.

18.4.1.1 ANSI92

ANSI92 = Yes|No

Values	Yes/No
Default	No
Description	Specifies whether the SQL generated complies to the ANSI92 standard. Yes: Enables the SQL generation compliant to ANSI92 standard.

No: SQL generation behaves according to the `PRM` parameter `OUTER_JOIN_GENERATION`.

18.4.1.2 AUTO_UPDATE_QUERY

`AUTO_UPDATE_QUERY` = Yes|No

Values	Yes/No
Default	No
Description	<p>Determines what happens when an object in a query is not available to a user profile.</p> <p>Yes: Query is updated and the object is removed from the query.</p> <p>No: Object is kept in the query.</p>

18.4.1.3 BEGIN_SQL

`BEGIN_SQL` = <String>

Values	String
Default	Empty string
Description	<p><code>BEGIN_SQL</code> is used to prefix SQL statements for accounting, prioritization, and workload management. The parameter applies to any SQL generation, including document generation and list of values queries.</p> <p><code>BEGIN_SQL</code> is supported in Web Intelligence , LiveOffice, Crystal Reports for Enterprise, and QaaWS. It is ignored by Desktop Intelligence.</p> <p>Example for Teradata:</p> <pre>BEGIN_SQL=SET QUERY_BAND='string' for transaction;</pre> <p>This parameter requires a string that contains one or more name-value pairs, separated by a semicolon, all inside single quotes. All SQL statements are prefixed with the parameter that follows <code>BEGIN_SQL</code>. The name-value pairs entered in this parameter are written in the <code>GetQueryBandPairs</code> system table.</p> <p>Example of three name-value pairs:</p> <pre>BEGIN_SQL=SET QUERY_BAND='UserID=Jones;JobID=980;AppID=TRM' for transaction;</pre>

You can also use the @Variable function as the value in the name-value pair, the returned value is enclosed in single quotes: `BEGIN_SQL=SET QUERY_BAND='USER=@Variable('BOUSER');Document=@Variable('DPNAME')';' for transaction;`

18.4.1.4 BLOB_COMPARISON

BLOB_COMPARISON = Yes|No

Values	Yes/No
Default	No
Can be edited?	No
Description	<p>Species if a query can be generated with a <code>DISTINCT</code> statement when a BLOB file is used in the <code>SELECT</code> statement. It is related to the setting <code>No Duplicate Row</code> in the query properties.</p> <p>Yes: The <code>DISTINCT</code> statement can be used within the query.</p> <p>No: The <code>DISTINCT</code> statement cannot be used within the query even if the query setting <code>No Duplicate Row</code> is on.</p>

18.4.1.5 BOUNDARY_WEIGHT_TABLE

BOUNDARY_WEIGHT_TABLE = Integer 32bits [0-9]

Values	Integer 32bits [0-9, or a negative integer]
Default	-1
Description	<p>Allows you to optimize the <code>FROM</code> clause when tables have many rows.</p> <p>If the table size (number of rows) is greater than the entered value, the table is declared as a subquery:</p> <pre>FROM (SELECT col1, col2,....., coln, ,....., FROM Table_Name WHERE simple condition).</pre> <p>A simple condition is defined as not having a subquery.</p> <p>-1, 0, or any negative number means that this optimization is not used.</p>
Limitations	<p>Optimization is not implemented when:</p> <ul style="list-style-type: none"> The operator <code>OR</code> is in the query condition

- Only one table is involved in the SQL
- The query contains an outer join
- No condition is defined on the table that is being optimized
- The table being optimized is a derived table.

18.4.1.6 CUMULATIVE_OBJECT_WHERE

CUMULATIVE_OBJECT_WHERE = Yes|No

Values	Yes No
Default	No
Description	<p>This parameter applies to filtered objects only. Specifies how to combine the objects WHERE clause with the query condition on those objects.</p> <p>Yes: Specifies that WHERE clauses are combined with the main query condition with the AND operator.</p> <p>No : Specifies that the object's WHERE clause is combined with the condition for this object.</p> <p>Example:</p> <p>If the condition is find all French clients different from John or American cities different from New York, the SQL is:</p> <p>Yes:</p> <pre>(customer.first_name <> 'John') OR (city.city <> 'New York AND customer_country.country = 'France' AND city_country.country = 'USA'</pre> <p>No:</p> <pre>(customer.first_name <> 'John' AND customer_country.country = 'France') OR (city.city <> 'New York' AND city_country.country = 'USA')</pre>

18.4.1.7 DISABLE_ARRAY_FETCH_SIZE_OPTIMIZATION

DISABLE_ARRAY_FETCH_SIZE_OPTIMIZATION = Yes|No

Values	Yes/No
Default	No
Description	<p>An optimization algorithm can be used to optimize the size of the returned arrays instead of using the default setting.</p> <p>No: All queries run on the universe will benefit from the optimization.</p> <p>Yes: Queries use the default value set.</p>

18.4.1.8 DISTINCT_VALUES

DISTINCT_VALUES = GROUPBY|DISTINCT

Values	GROUPBY DISTINCT
Default	DISTINCT
Description	<p>Specifies whether SQL is generated with a <code>DISTINCT</code> or <code>GROUP BY</code> clause for objects in the business layer, and in lists of values. In the Query Panel, query generation takes the value of <code>DISTINCT_VALUES</code> into account only when the option Retrieve duplicate rows is unselected in the query properties.</p> <p><code>DISTINCT</code>: The SQL is generated with a <code>DISTINCT</code> clause, for example:</p> <pre>SELECT DISTINCT cust_name FROM Customer</pre> <p><code>GROUPBY</code>: The SQL is generated with a <code>GROUP BY</code> clause, for example:</p> <pre>SELECT cust_name FROM Customer GROUP BY Customer.cust_name</pre>

18.4.1.9 END_SQL

END_SQL = String

Values	String
Default	<empty string>
Description	The statement specified in this parameter is added at the end of each SQL statement.
Example	<p>For IBM DB2 databases, you can use the following:</p> <pre>END_SQL=FOR SELECT ONLY</pre> <p>The server will read blocks of data much faster.</p>

	<p>Another example:</p> <pre>END_SQL='write ` UNVID To Usage_Audit.Querieded_universe</pre> <p>Would write universe id to an audit table, this can be used to record other data such as user and tables queried.</p>
--	--

18.4.1.10 EVAL_WITHOUT_PARENTHESIS

EVAL_WITHOUT_PARENTHESIS = Yes|No

Values	Yes No
Default	No
Description	<p>By default, the function @Select(folder\object) is replaced by the SELECT statement for the object <folder\object> enclosed within brackets.</p> <p>For example, when combining two @Select statements, @Select(object1) * @Select(object2).</p> <p>If the SQL(object1) = A-B and SQL(object2) = C,</p> <p>then the operation is (A-B) * (C).</p> <p>You avoid the default adding of brackets by setting EVAL_WITHOUT_PARENTHESIS = Yes. The operation is then A - B * C.</p> <p>Yes: Brackets are removed from the SELECT statement for a function @Select(folder\object)</p> <p>No: Brackets are added around the Select statement for the function @Select(folder\object).</p>

18.4.1.11 FILTER_IN_FROM

FILTER_IN_FROM = Yes|No

Values	Yes / No
Default	No
Description	Determines if the generated SQL includes query filters in the FROM clause whenever possible.

i Note

This setting is only applicable if the SQL generation parameter `ANSI92` is set to `Yes`.

This parameter is useful when querying tables that have outer joins defined. For example, an outer join on tables `Customer` and `Reservations` returns all customers, even those without reservations. A query filter in the `WHERE` clause might filter out customers without reservations. If the `FILTER_IN_FROM` parameter is set to `Yes`, the generated SQL includes query filters in the `FROM` clause whenever possible in order to preserve the null values returned by the outer join.

`Yes`: When SQL is generated, query filters are put in the `FROM` clause whenever possible.

`No`: When SQL is generated, query filters are put in the `WHERE` clause.

18.4.1.12 FORCE_SORTED_LOV

`FORCE_SORTED_LOV` = Yes|No

Values	Yes No
Default	No
Description	Retrieves a list of values that is sorted. <code>Yes</code> : Specifies that the list of values is sorted. <code>No</code> : Specifies that the list of values is not sorted.

18.4.1.13 GROUPBY_PRIMARY_KEY

`GROUPBY_PRIMARY_KEY` = YES | NO

Values	YES NO
Default	YES
Description	Allows you to deactivate the use of the primary key in the <code>GROUP BY</code> clause. By default, if data for an index aware object is retrieved, the SQL is optimized by using the primary key in the <code>GROUP BY</code> clause. <code>YES</code> : Favors the use of the primary key over the column name in the <code>GROUP BY</code> clause.

NO: Does not use the primary key in the GROUP BY clause.

18.4.1.14 INNERJOIN_IN_WHERE

INNERJOIN_IN_WHERE = Yes|No

Values	Yes No
Default	No. You must manually add the parameter to activate it.
Description	<p>Allows you to force the system to generate SQL syntax with all the inner joins in the WHERE clause when ANSI92 is set to yes. This is only possible if a query contains only inner joins (Does not contain FULL OUTER, RIGHT OUTER, or LEFT OUTER joins).</p> <p>Yes: If ANSI92 is set to yes, the system generates ANSI92 join syntax in the FROM clause except when the query contains only inner joins. In this case, the inner joins go into the WHERE clause.</p> <p>No: If ANSI92 is set to Yes, the system generates ANSI 92 join syntax in the FROM clause.</p>

18.4.1.15 JOIN_BY_SQL

JOIN_BY_SQL = Yes|No

Values	Yes No
Default	No
Description	<p>Specifies how multiple SQL statements are handled. Multiple statements can be combined (provided that the database permits this).</p> <p>Yes: Specifies that multiple SQL statements are combined.</p> <p>No: Specifies that multiple SQL statements are not combined. This is the default value.</p>

18.4.1.16 MAX_INLIST_VALUES

MAX_INLIST_VALUES = [0-99]

Values	Integer: min-1, max depends on DB
Default	-1
Description	<p>Allows you to set the maximum number of values you may enter in a condition when you use the <code>IN LIST</code> operator.</p> <p>99: Specifies that you may enter up to 99 values when you create a condition using the <code>IN LIST</code> operator.</p> <p>The maximum authorized value you may enter depends on your database.</p> <p>The value of -1 means that there is no restriction on the number of values returned, except that imposed by the database.</p>

18.4.1.17 REPLACE_COMMA_BY_CONCAT

REPLACE_COMMA_BY_CONCAT= Yes|No

Values	Yes No
Default	No
Description	<p>In previous versions of the universe design tool, a comma could be used to separate multiple fields in an object Select statement. The comma was treated as a concatenation operator. For universes that already use the comma in this way you can set <code>REPLACE_COMMA_BY_CONCAT</code> to <code>No</code> to keep this behavior. In the current version of the universe design tool, this parameter is set to <code>Yes</code> by default, so that any expressions using a comma in this way are automatically changed to use concatenation syntax.</p> <p><code>Yes</code>: Comma is replaced by the concatenation expression when multi field object is found.</p> <p><code>No</code>: Keep the comma as it is.</p>

18.4.1.18 SELFJOINS_IN_WHERE

SELFJOINS_IN_WHERE = Yes|No

Values	Yes No
Default	No
Description	<p>Self-joins are usually included in the FROM clause. This allows you to force the system to generate SQL syntax with all the conditions of a self-join in the <code>WHERE</code></p>

clause. The `ANSI92` parameter must be set to `Yes` for this parameter to be taken into account.

You must manually add the parameter to the list to activate it.

`Yes`: The conditions of a self-join go in the `WHERE` clause of the SQL query.

`No`: The syntax for self-joins is generated according to the ANSI 92 convention, and conditions for a self-join go in the `ON` clause of the table join definition in the `FROM` clause of the SQL query.

18.4.1.19 SHORTCUT_BEHAVIOR

`SHORTCUT_BEHAVIOR` = `ShortestPath`|`Global`|`Successive`

Values	<code>ShortestPath</code> <code>Global</code> <code>Successive</code>
Default	<code>ShortestPath</code>
Description	<p>Specifies how shortcut joins are applied.</p> <p><code>ShortestPath</code>: applies shortcuts so as to obtain the smallest number of tables in the query.</p> <p><code>Successive</code>: applies shortcuts one after the other. If a shortcut removes a table involved in a potential successive shortcut, the successive shortcut is not applied.</p> <p><code>Global</code>: applies all shortcuts. If the resulting query creates a Cartesian product, no shortcut joins are applied.</p> <div> <p>i Note</p> <p>This parameter was formerly listed as <code>GLOBAL_SHORTCUTS</code> in the <code>PRM</code> files. The value <code>Global</code> corresponds to <code>Yes</code>, and <code>Successive</code> corresponds to <code>No</code>.</p> </div>

18.4.1.20 SMART_AGGREGATE

`SMART_AGGREGATE` = `Yes`|`No`

Values	<code>Yes</code> <code>No</code>
Default	<code>No</code>
Description	Determines how aggregate tables are used for smart measures that are based on an aggregate tables. This ensures that a universe object based on a ratio is

	<p>correctly aggregated. By default the system takes the advantage of the pre-calculated values from the aggregated tables, if these table are not consistent during time (different time periods), you use this parameter to ensure the most detailed aggregate tables are used.</p> <p>This parameter is not visible in the universe parameter list (by default not activated). The universe designer must manually insert it in the parameter list before activating it (value <code>Yes</code>).</p> <p><code>Yes</code>: Any additional grouping set query should be based on the aggregate table of the initial query for the smart measure based on aggregate table.</p> <p><code>No</code>: The system takes the most appropriate aggregate table.</p>
--	--

18.4.1.21 THROUGH_AGGREGATE_AWARE

THROUGH_AGGREGATE_AWARE = Yes|No

Values	Yes/No
Default	<p>Yes</p> <div> <p>i Note</p> <p>For universes converted from .unv, the default value is No.</p> </div>
Description	<p>Determines if aggregate awareness is taken into account when testing the compatibility of query objects.</p> <p>This parameter allows you to possibly improve the results of queries on converted .unv universes that fail for the .unv universe.</p> <p><code>Yes</code>: Aggregate awareness is taken into account when testing the compatibility of objects in the query. In some cases this allows the query to succeed in situations where there are incompatible objects (split queries) with aggregate aware objects.</p> <p><code>No</code>: Compatibility testing of objects uses the behavior for .unv universes.</p>

18.4.1.22 THOROUGH_PARSE

THOROUGH_PARSE = Yes|No

Values	Yes No
Default	No

Description	<p>Specifies the methodology used for default Parsing in the Query pane and individual object parsing.</p> <p>Yes: PREPARE, DESCRIBE, and EXECUTE statements are used to parse SQL for objects.</p> <p>Prepare+DescribeCol+Execute</p> <p>No: PREPARE and DESCRIBE statements are used to parse SQL for objects.</p>
-------------	--

18.4.1.23 TRUST_CARDINALITIES

TRUST_CARDINALITIES = Yes|No

Values	Yes No
Default	No
Description	<p>Allows you to optimize the SQL in case of inflated results.</p> <p>Yes: For queries that include a measure, all conditions that inflate the measure and do not appear in the Result Objects, are transformed to sub queries to ensure that tables that may return false results for the measure are not included in the query.</p> <p>No: No optimization is implemented.</p>

18.4.1.24 UNICODE_STRINGS

UNICODE_STRINGS = Yes|No

Values	Yes No
Default	No
Description	<p>Specifies whether the current universe can manipulate Unicode strings or not. Only applies to Microsoft SQL Server and Oracle 9. If the database character set in the SBO file is set as Unicode, then it is necessary to modify the SQL generation to handle specific Unicode column types like NCHAR and NVARCHAR.</p> <p>Yes: Conditions based on strings are formatted in the SQL according to the value for a parameter UNICODE_PATTERN in the PRM file, for example for MS SQL Server (sqlsrv.prm) : UNICODE_PATTERN=N\$</p> <p>The condition Customer_name='Arai ' becomes</p> <p>Customer_name=N'Arai'.</p>

	<p>Note: When you create a prompt with @Prompt syntax based on Unicode value, the data type should be 'U' not 'C'</p> <p>No: All conditions based on strings are formatted in the standard SQL. For example the condition Customer_name='Arai ' remains Customer_name='Arai'</p>
--	--

18.4.1.25 USE_ENHANCED_QUERY_STRIPPING

USE_ENHANCED_QUERY_STRIPPING = Yes|No

Values	Yes No
Default	No
Description	<p>Specifies the query stripping mode for relational universes. If set to Yes, the system optimizes only the SELECT and GROUP BY clauses to avoid fetching unused data but does not modify the other clauses to respect the original query semantic.</p> <p>If set to No or not set, the system generates optimized queries by completely ignoring the stripped objects with their corresponding tables and joins.</p> <div style="background-color: #fff9c4; padding: 10px; margin: 10px 0;"> <p>i Note</p> <p>If aggregate awareness is defined in the business layer (using the @Aggregate_aware function in the definition of business layer objects), enhanced query stripping is used regardless of the value of USE_ENHANCED_QUERY_STRIPPING.</p> </div> <p>For more information, see About query stripping [page 229].</p>

18.4.2 SQL generation parameters set in the extended PRM

The following reference describes the SQL generation parameters that you set in the extended data access parameter (PRM) file for the target data access driver. Extended PRM files are located in the the following directory, where <RDBMS> is the network layer or middleware name:

```
<BIP_INSTALL_DIR>\SAP BusinessObjects Enterprise XI 4.0\dataAccess\connectionServer
\<RDBMS>\extensions\qt
```

For more information about PRM files, see the *Data Access Guide*.

Related Information

[CASE_SENSITIVE](#) [page 446]

[COMMA](#) [page 446]

[CONCAT](#) [page 447]

[DELIMIT_IDENTIFIERS](#) [page 447]

[DELIMIT_LOWERCASE](#) [page 447]

[EXTERN_SORT_EXCLUDE_DISTINCT](#) [page 448]

[GROUPBY_WITH_ALIAS](#) [page 448]

[IDENTIFIER_DELIMITER](#) [page 448]

[OUTERJOINS_GENERATION](#) [page 449]

[OVER_CLAUSE](#) [page 451]

[OWNER](#) [page 451]

[QUALIFIER](#) [page 451]

[UNICODE_PATTERN](#) [page 451]

[USER_INPUT_DATE_FORMAT](#) [page 452]

[USER_INPUT_NUMERIC_SEPARATOR](#) [page 452]

18.4.2.1 CASE_SENSITIVE

```
<Parameter Name="CASE_SENSITIVE">NO</Parameter>
```

Description	Specifies if the database is case-sensitive. This parameter is used with Oracle.
Values	YES: the database is case-sensitive. NO: the database is not case-sensitive.
Default	NO

18.4.2.2 COMMA

```
<Parameter Name="COMMA">|| ' ' ||</Parameter>
```

Description	Specifies what database concatenation operator should be used to replace a comma for objects that have the following syntax: Tab.Col1, Tab.Col2. This parameter is used with all data access drivers.
Values	' ' + ' '+

Default	' '
Result	Tab.Col1 ' ' Tab.Col2

18.4.2.3 CONCAT

<Parameter Name="CONCAT">| |</Parameter>

Description	Specifies the concatenation operator. The parameter is used with all data access drivers.
Values	double pipe () or plus sign (+)
Default	

18.4.2.4 DELIMIT_IDENTIFIERS

<Parameter Name="DELIMIT_IDENTIFIERS">YES</Parameter>

Description	Specifies if database identifiers can be quoted. Identifiers are quoted using the delimiter specified in the IDENTIFIER_DELIMITER parameter.
Values	YES: identifiers can be quoted. NO: identifiers cannot be quoted.
Default	YES
Result	Table name="my_table"

18.4.2.5 DELIMIT_LOWERCASE

<Parameter Name="DELIMIT_LOWERCASE"></Parameter>

Description	Specifies if lowercase identifiers are delimited with quotes.
Values	YES: the lowercase identifiers are delimited with quotes. NO: the lowercase identifiers are not delimited with quotes.

18.4.2.6 EXTERN_SORT_EXCLUDE_DISTINCT

<Parameter Name="EXTERN_SORT_EXCLUDE_DISTINCT">YES</Parameter>

Description	Specifies if the application generates a SELECT DISTINCT when a query contains an ORDER BY clause.
Values	YES: a SELECT DISTINCT is not generated when the query contains an ORDER BY clause. NO: a DISTINCT is generated when the query contains an ORDER BY clause.
Default	YES

18.4.2.7 GROUPBY_WITH_ALIAS

<Parameter Name="GROUPBY_WITH_ALIAS">YES</Parameter>

Description	Specifies if the database can create a GROUP BY clause that contains aliases in the SELECT statement.
Values	YES: it allows you to create a GROUP BY clause with aliases in the SELECT statement. NO: it does not let you create a GROUP BY clause with aliases in the SELECT statement.
Default	YES

18.4.2.8 IDENTIFIER_DELIMITER

<Parameter Name="IDENTIFIER_DELIMITER">"</Parameter>

Description	<p>Specifies the following features:</p> <ul style="list-style-type: none">• Table or column names that contain spaces or special characters are enclosed within quotes if BACK_QUOTE_SUPPORTED parameter is activated.• Tables or column names regardless of their characters are enclosed within quotes if DELIMIT_IDENTIFIERS parameter is activated. <p>To use this parameter, either BACK_QUOTE_SUPPORTED or DELIMIT_IDENTIFIERS must be set to YES. This is the default value of both parameters.</p>
Values	" (double quote): table or column names that contain spaces or special characters are enclosed in double quotes.

	' (single quote): table or column names that contain spaces or special characters are enclosed in single quotes. This value can only be used with Microsoft Access.
Default	"
Result	Table name="My Table"

18.4.2.9 OUTERJOINS_GENERATION

<Parameter Name="OUTERJOINS_GENERATION">ANSI92</Parameter>

Description	<p>Specifies the SQL syntax for outer joins.</p> <p>The value <code>ANSI_92</code> generates an outer join in the FROM clause. Other values generate the outer join in the WHERE clause.</p> <p>When you modify this setting, you should check join properties to verify that the outer join expression is valid, and that the cardinalities are correct. ANSI92 does not support any manual customization in the join syntax.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>The PRM file <code>OUTERJOINS_GENERATION</code> parameter relates to the universe <code>ANSI92</code> setting in the following way:</p> <ul style="list-style-type: none"> • If the PRM file <code>OUTERJOINS_GENERATION</code> parameter is set to <code>ANSI_92</code> and the universe <code>ANSI92</code> parameter is set to <code>NO</code>, the PRM parameter overrides the universe setting and outer joins conform to ANSI92 behavior. • If the PRM file <code>OUTERJOINS_GENERATION</code> parameter is set to <code>USUAL</code>, then the universe <code>ANSI92</code> setting takes precedence, and outer joins conform to ANSI92 depending on whether the universe <code>ANSI92</code> setting is <code>YES</code> or <code>NO</code>. </div>
Values	<p>The primary values for <code>OUTERJOINS_GENERATION</code> are:</p> <ul style="list-style-type: none"> • <code>ANSI_92</code>: the default outer join behavior conforms to the ANSI92 standard regardless of the <code>ANSI92</code> parameter value in the universe. • <code>NO</code>: outer joins are not supported. • <code>USUAL</code>: the database-specific outer join behavior is used. This behavior is overridden if the <code>ANSI92</code> parameter is set to <code>YES</code>. <p>Other settings are available depending on the database. See the defaults below.</p>
Default	<p><code>ANSI_92</code>: default value for Oracle, MS SQL Server 2005 and Sybase.</p> <p><code>DB2</code>: default value for IBM DB2.</p>

<p>FULL_ODBC: default value for Microsoft SQL Server.</p> <p>INFORMIX: default value for IBM Informix.</p> <p>INGRES: default value for Teradata.</p> <p>NO: default value for ODBC.</p> <p>USUAL: default value for HP Neoview, Netezza, IBM Red Brick and MS SQL Server 2000.</p>

Examples of OUTERJOINS_GENERATION parameter settings

Setting = USUAL:

```
FROM T1, T2
WHERE T1.col1(+) = T2.col2
```

Setting = DB2:

```
FROM T2 LEFT OUTER JOIN T1
ON T1.col1 = T2.col2
```

Setting = ODBC:

```
FROM {oj T1 LEFT OUTER JOIN T2 ON T1.col1=T2.col2}
Where (T2.col3 = T3.col1)
```

Setting = INFORMIX

```
FROM T2
OUTER T1
WHERE T1.col1=T2.col2
```

Setting = FULL-ODBC

```
FROM {oj T1 RIGHT OUTER JOIN T2 ON T2.col2=T1.col1
T2 INNER JOIN 3 on T2.col3 = T3.col1}
```

Setting = ANSI_92:

```
SELECT DISTINCT
  t1.col1,
  t2.col2
FROM
  (t1 RIGHT OUTER JOIN t2 ON (t1.col1=t2.col2) )
```

18.4.2.10 OVER_CLAUSE

<Parameter Name="OVER_CLAUSE">YES</Parameter>

Description	Allows SAP BusinessObjects applications to include RISQL functions when generating SQL. The supported RISQL functions for the database are listed in the ANALYTIC_FUNCTIONS parameter.
Values	YES: applications can include RISQL functions when generating SQL. NO: applications cannot include RISQL functions when generating SQL.
Default	YES

18.4.2.11 OWNER

<Parameter Name="OWNER">YES</Parameter>

Description	Specifies if the database supports the owner name as prefix for tables.
Values	YES: the database supports prefixing tables with the owner name. NO: the database does not support prefixing tables with the owner name.
Default	YES

18.4.2.12 QUALIFIER

<Parameter Name="QUALIFIER">NO</Parameter>

Description	Specifies if the database supports the qualifier name as prefix for tables.
Values	YES: the database supports prefixing tables with the qualifier name. NO: the database does not support prefixing tables with the qualifier name.
Default	RDBMS-dependent.

18.4.2.13 UNICODE_PATTERN

<Parameter Name="UNICODE_PATTERN">UNISTR(\$)</Parameter>

Description	Only applies when the universe SQL generation parameter <code>UNICODE_STRINGS</code> is set to <code>YES</code> . All conditions based on strings are then formatted with this string value. This is used with MS SQL Server and Oracle only.
Values	<code>N\$</code> : for MS SQL Server <code>UNISTR (\$)</code> : for Oracle

18.4.2.14 USER_INPUT_DATE_FORMAT

```
<Parameter Name="USER_INPUT_DATE_FORMAT">'dd-MM-yyyy HH:mm:ss'</Parameter>
```

Description	Specifies the default date and hour formats generated in the WHERE clause of a SQL statement.
Values	<p><code>{\d 'yyyy-mm-dd'}</code>: default date format with ODBC.</p> <p><code>'DD-MM-YYYY HH:MM:SS'</code>: default date and hour formats with Oracle.</p> <p><code>'MM/DD/YYYY'</code>: default date format with IBM Informix.</p> <p><code>'yyyy-mm-dd HH:mm:ss'</code>: default date and hour formats with MS SQL Server and for most IBM DB2 servers.</p> <p><code>'mm/dd/yyyy hh:m:s am/pm'</code>: default date and hour formats with Sybase.</p> <p><code>'yyyy-mm-dd'</code>: default date format with a Sybase gateway.</p> <div style="background-color: #fff9c4; padding: 10px; margin-top: 10px;"> <p>i Note</p> <p>If you need to use time or timestamp variables with ODBC, you must replace the default date format value with <code>{\t 'hh:mm:ss'}</code> or <code>{\t\s 'yyyy-mm-dd hh:mm:ss'}</code> in the <code>odbc.sbo</code> file.</p> </div>
Default	See values above.

18.4.2.15 USER_INPUT_NUMERIC_SEPARATOR

```
<Parameter Name="USER_INPUT_NUMERIC_SEPARATOR">.</Parameter>
```

Description	Specifies the default decimal separator that is used in the generated SQL script.
Values	<code>'.'</code> (period)
Default	<code>'.'</code>



www.sap.com/contactsap

© 2014 SAP AG or an SAP affiliate company. All rights reserved.
No part of this publication may be reproduced or transmitted in any form or for any purpose without the express permission of SAP AG. The information contained herein may be changed without prior notice.

Some software products marketed by SAP AG and its distributors contain proprietary software components of other software vendors. National product specifications may vary.

These materials are provided by SAP AG and its affiliated companies ("SAP Group") for informational purposes only, without representation or warranty of any kind, and SAP Group shall not be liable for errors or omissions with respect to the materials. The only warranties for SAP Group products and services are those that are set forth in the express warranty statements accompanying such products and services, if any. Nothing herein should be construed as constituting an additional warranty.

SAP and other SAP products and services mentioned herein as well as their respective logos are trademarks or registered trademarks of SAP AG in Germany and other countries.

Please see <http://www.sap.com/corporate-en/legal/copyright/index.epx> for additional trademark information and notices.