



MASHZONE NEXTGEN

USER AND DEVELOPER GUIDE

VERSION 10.3

October 2018

This document applies to MashZone NextGen Version 10.3 and to all subsequent releases. Specifications contained herein are subject to change and these changes will be reported in subsequent release notes or new editions.

Copyright © 2013 - 2018 Software AG, Darmstadt, Germany and/or Software AG USA Inc., Reston, VA, USA, and/or its subsidiaries and/or its affiliates and/or their licensors.

The name Software AG and all Software AG product names are either trademarks or registered trademarks of Software AG and/or Software AG USA Inc. and/or its subsidiaries and/or its affiliates and/or their licensors. Other company and product names mentioned herein may be trademarks of their respective owners.

Detailed information on trademarks and patents owned by Software AG and/or its subsidiaries is located at <http://softwareag.com/licenses>.

Use of this software is subject to adherence to Software AG's licensing conditions and terms. These terms are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

This software may include portions of third-party products. For third-party copyright notices, license terms, additional rights or restrictions, please refer to "License Texts, Copyright Notices and Disclaimers of Third Party Products". For certain specific third-party license restrictions, please refer to section E of the Legal Notices available under "License Terms and Conditions for Use of Software AG Products / Copyright and Trademark Notices of Software AG Products". These documents are part of the product documentation, located at <http://softwareag.com/licenses> and/or in the root installation directory of the licensed product(s).

Contents

1	Preface.....	1
2	Introducing Software AG MashZone NextGen	2
3	MashZone NextGen welcome page.....	3
4	Open a dashboard in MashZone NextGen.....	4
5	Use the dashboard editor.....	5
5.1	Use dashboards in view mode	5
5.1.1	Display a dashboard in view mode	5
5.1.2	Use interactive filters in dashboards	5
5.1.3	Use the filter panel of Process Mining context-based widgets.....	6
5.1.4	Refresh data of widgets	7
5.1.5	Pause automatic data refreshing.....	8
5.1.6	Change column width and sort order	8
5.1.7	Use lists for multiple selection	9
5.1.8	Save widget data as a CSV file	10
5.2	Manage dashboards	10
5.2.1	Open a dashboard in the dashboard editor	10
5.2.2	Create a dashboard in the dashboard editor	11
5.2.3	Use smart dashboard edit mode	11
5.2.3.1	Insert new row	12
5.2.3.2	Delete a row.....	12
5.2.3.3	Resize a row.....	12
5.2.3.4	Lock row height	12
5.2.3.5	Change a row order.....	13
5.2.3.6	Place a widget	13
5.2.3.7	Resize widgets.....	13
5.2.3.8	Avoid line break in dashboard view mode.....	13
5.2.3.9	Group widgets	14
5.2.3.10	Set row style.....	14
5.2.3.11	Set layout group style	14
5.2.3.12	Set dashboard tab style.....	15
5.2.3.13	Switch to fixed-grid workspace	15
5.2.4	Edit a dashboard.....	15
5.2.5	Delete a dashboard	16
5.2.6	Edit dashboard properties	16
5.2.7	Manage dashboard permissions	16
5.2.8	Change dashboard style template	18
5.2.9	Use the Process Mining context.....	18
5.2.9.1	Create a Process Mining context	18
5.2.9.2	Edit a Process Mining context	20
5.2.9.3	Delete a Process Mining context	21
5.2.9.4	Assign data columns to Process Mining context-based widgets.....	21
5.2.9.5	Configure filters for Process Mining context-based widgets.....	22
5.2.9.5.1	Configure a date filter for a Process Mining context	23
5.2.9.6	Exclude a widget from a Process Mining context filter	24
5.2.9.7	Enable the filter panel.....	24
5.2.9.8	Configure the filter panel	24
5.2.9.9	Migrate a dashboard context	26
5.2.10	Use Multi-lingual dashboards.....	27
5.2.10.1	Configure languages for translation	27

5.2.10.2	Translate the widget content	27
5.2.10.3	Select the language for your dashboard	28
5.2.10.4	Translate the tab title	28
5.3	Configure widgets	29
5.3.1	Set up widgets	30
5.3.2	Insert widgets in a dashboard	30
5.3.3	Copy and paste widgets in dashboards	31
5.3.4	Assign data sources to widgets	31
5.3.4.1	Assign data columns to line, column or bar charts	32
5.3.4.1.1	Set properties of widget elements	33
5.3.4.1.2	Configure KPI thresholds	34
5.3.4.1.3	Configure KPI rating	35
5.3.4.2	Assign data columns to pie charts	36
5.3.4.2.1	Set properties of widget elements	36
5.3.4.3	Assign data columns to bubble charts	37
5.3.4.3.1	Set properties of widget elements	38
5.3.4.3.2	Configure KPI thresholds	39
5.3.4.4	Assign data columns to grids	40
5.3.4.4.1	Set properties of widget elements	40
5.3.4.4.2	Configure thresholds	41
5.3.4.5	Assign data columns to lists	43
5.3.4.5.1	Set properties of widget elements	43
5.3.4.6	Assign data columns to circular gauge charts	44
5.3.4.6.1	Set properties of widget elements	45
5.3.4.6.2	Configure thresholds	45
5.3.4.7	Assign data columns to horizontal and vertical gauge charts	46
5.3.4.7.1	Set properties of widget elements	46
5.3.4.7.2	Configure thresholds	47
5.3.4.8	Assign data columns to traffic lights	48
5.3.4.8.1	Set properties of widget elements	48
5.3.4.8.2	Configure thresholds	49
5.3.4.9	Assign data columns to a drop-down box	49
5.3.4.9.1	Set properties of widget elements	50
5.3.4.10	Assign data columns to labels	51
5.3.4.10.1	Set properties of widget elements	51
5.3.4.11	Assign data columns to a rich text area	52
5.3.4.11.1	Set properties of widget elements	52
5.3.4.12	Assign data columns to sliders	53
5.3.4.12.1	Set properties of widget elements	53
5.3.4.13	Assign data columns to date filters	54
5.3.4.14	Assign data columns to vector map	54
5.3.4.14.1	Set properties of widget elements	55
5.3.4.14.2	Configure KPI thresholds	56
5.3.4.15	Assign data columns to map with markers	56
5.3.4.15.1	Set properties of widget elements	57
5.3.4.15.2	Configure KPI thresholds	58
5.3.4.16	Create input parameters	59
5.3.5	Display legend	60
5.3.6	Remove widgets from dashboard	60

5.3.7	Define filters for widgets	60
5.3.8	Specify actions for widgets.....	62
5.3.8.1	Change tab	62
5.3.8.2	Set selection	63
5.3.8.3	Call URL.....	64
5.3.8.4	Specify an image source URL.....	65
5.3.8.5	Post data	66
5.3.8.6	Send events	67
5.3.9	Hide or display widget header and border	68
5.3.10	Resize widgets.....	68
5.3.11	Move widgets to the foreground or background	69
5.3.12	Display widgets on multiple dashboard tabs	69
5.3.13	Change widget style template.....	70
5.4	Configure dashboard tabs	70
5.4.1	Add dashboard tabs	70
5.4.2	Delete dashboard tabs.....	70
5.4.3	Set dashboard tab properties	71
5.4.4	Set dashboard tab style.....	71
5.4.5	Set dashboard tab as default.....	72
5.5	Other	72
5.5.1	Use dynamic URL selection.....	72
5.5.2	Hide tab bar in view mode	73
5.5.3	Set display size	74
5.5.4	Set grid lines.....	74
5.5.5	Allow anonymous access to dashboards	75
5.5.5.1	PPM data source operator	75
5.5.5.2	File resources	75
5.5.6	Customize the MashZone NextGen welcome page	76
5.5.7	Display dashboards without application header	77
5.5.8	Export dashboards	77
5.5.9	Import dashboards.....	78
5.5.10	Create widget style template.....	79
6	Use the data feed editor	81
6.1	Create a data feed	81
6.1.1	Open a new data feed	81
6.1.2	Select data sources	81
6.1.3	Calculate the feed data.....	82
6.1.4	Add further operators	82
6.1.5	Connect the operators	83
6.1.6	Save the data feed.....	83
6.2	Edit a data feed	83
6.3	Delete data feeds.....	84
6.4	Edit data feed properties	84
6.5	Manage data feed permissions.....	84
6.6	Change the feed editor style	86
6.7	Data feed operators	86
6.8	Use Materialized Feeds	86
6.8.1	Use datasets storing data feed results	87
6.8.2	Administrate Materialized Feeds in Terracotta DB.....	88
6.8.3	Use datasets from external application	88
6.8.4	Which privileges are required?.....	89
6.9	Other	89
6.9.1	Import data feeds	90

6.9.2	Export data feeds	90
6.9.3	Display data feed without application header	92
7	Valuable information	93
7.1	What are dashboards?.....	93
7.2	What are data feeds?	93
7.3	What is the view-only license?.....	94
8	Administration.....	95
8.1	Open the Admin Console	95
8.2	MashZone NextGen Security	95
8.2.1	Manage your MashZone NextGen profile	96
8.2.1.1	Manage your locale and account information	96
8.2.1.2	Change your password	96
8.2.2	MashZone NextGen server	97
8.2.3	MashZone NextGen repository	97
8.2.4	Change technical user password	97
8.2.5	Authentication and Guest Access	98
8.2.5.1	User Authentication	98
8.2.5.2	Valid Credentials.....	99
8.2.5.3	Sessions and Timeouts	99
8.2.6	Default User Accounts	100
8.2.7	Authentication with Single Sign-On Solutions	100
8.2.7.1	Configuration for Agent-Based SSO Solutions	101
8.2.7.2	Implementing a Custom SSO Filter.....	103
8.2.7.3	SSO integration in My webMethods	103
8.2.8	Authorization Policies and Permissions.....	104
8.2.8.1	Grant User Access to MashZone NextGen with Built-in Groups	104
8.2.9	Built-In MashZone NextGen User Groups and Permissions.....	105
8.2.9.1	Access Policies Using MashZone NextGen Built-In Groups.....	106
8.2.10	Protect MashZone NextGen Event Service access	106
8.2.11	Anti-Clickjacking prevention when using iFrame.....	107
8.2.11.1	MashZone NextGen HTTP header security filter	107
8.2.11.1.1	Example.....	108
8.2.11.2	MashZone NextGen Content Security Policy.....	109
8.2.11.3	Add a trusted site to allow iFrame	109
8.2.11.4	Add multiple trusted sites to allow iFrame	110
8.2.11.5	Content-Security-Policy using wildcards	110
8.2.12	Handle personal data in log files	110
8.3	Getting Started with the MashZone NextGen Server	111
8.3.1	Additional MashZone NextGen System and Software Requirements.....	112
8.3.1.1	Additional Recommendations for MashZone NextGen.....	112
8.3.2	What is Installed with MashZone NextGen.....	113
8.3.3	Start and Stop the MashZone NextGen Server	113
8.3.3.1	Start the MashZone NextGen Event Service.....	113
8.3.3.2	Start the MashZone NextGen Server	114
8.3.3.3	Stop the MashZone NextGen Event Service	114
8.3.3.4	Stop the MashZone NextGen Server	115
8.3.4	Startup Considerations	115
8.3.5	Manage Licenses for MashZone NextGen and BigMemory	115
8.3.6	Move the MashZone NextGen repository to a robust database solution	116

8.3.6.1	Troubleshooting Connections to the MashZone NextGen Repository	117
8.3.6.2	Move the MashZone NextGen repository to Microsoft SQL Server	118
8.3.6.3	Move the MashZone NextGen repository to MySQL	120
8.3.6.4	Move the MashZone NextGen repository to Oracle.....	122
8.3.6.5	Move the MashZone NextGen repository to PostGres	124
8.3.7	Integrate Your LDAP Directory with MashZone NextGen.....	126
8.3.7.1	Defining LDAP Connection Configuration	128
8.3.7.2	Defining the Authentication Scheme	128
8.3.7.3	Defining the Authorization Scheme.....	129
8.3.7.4	Enabling MashZone NextGen Application Queries for All LDAP Users or Groups for Permissions	131
8.3.8	Use the Default MashZone NextGen User Repository	132
8.3.8.1	Manage Users.....	133
8.3.8.1.1	Create Users.....	133
8.3.8.1.2	Edit, Grant Permissions and other User Management Tasks ...	134
8.3.8.2	Manage User Groups	134
8.3.8.3	Automatically Assign New Users to Groups	135
8.3.8.4	Grant dashboard and data feed permissions via API console	136
8.3.8.4.1	Enable dashboard and data feed creation	136
8.3.8.5	Install MashZone NextGen and MashZone NextGen Event Service as Windows services.....	137
8.3.9	Command Central	137
8.3.9.1	MashZone NextGen plug-in	138
8.3.9.1.1	Required installer modules	138
8.3.9.1.2	Instance Overview	138
8.3.9.1.3	Instance Configuration.....	139
8.3.9.1.4	Instance Logs	141
8.3.9.1.5	Enable remote JMX connection.....	141
8.3.9.2	MashZone NextGen Event Service plug-in	141
8.3.9.2.1	Instance overview	141
8.3.9.2.2	Instance configuration	141
8.3.9.2.3	Instance Logs	144
8.3.9.3	Configure MashZone NextGen Digital Event Service (DES)	144
8.3.9.4	Configure ports.....	145
8.3.9.5	Configure keystores	146
8.3.9.6	Configure truststores.....	146
8.3.9.7	Configure SSL certificates	147
8.3.9.8	Manage licenses.....	147
8.4	MashZone NextGen Server Configuration	148
8.4.1	Memory Configuration for the MashZone NextGen Server	148
8.4.1.1	Configuration When MashZone NextGen Uses Only Heap Memory ...	148
8.4.1.2	Configuration When MashZone NextGen Uses Heap and Off-Heap Memory	149
8.4.2	Support International Character Sets and Locales	150
8.4.2.1	Set the Repository Character Encoding	151
8.4.2.2	Set the Repository Timezone or Offset	151
8.4.2.3	Date, Time and Numeric Display Options	152
8.4.2.4	Message Log and Default Locales	152
8.4.3	Edit style templates	152
8.4.4	Configure the MashZone NextGen server with custom ports	153
8.4.4.1	Change MashZone NextGen Server Ports	154

8.4.4.2	Change MashZone NextGen Repository Ports.....	154
8.4.4.3	Tomcat Application Server Port.....	154
8.4.5	Configure the MashZone NextGen server to work with a proxy server.....	155
8.4.6	Embedding MashZone NextGen in external system environments	155
8.4.6.1	Configure MashZone NextGen server to work with iFrame	156
8.4.6.2	Post data	157
8.4.6.3	URL selection.....	157
8.4.7	Define a Proxy Server Whitelist for MashZone NextGen	157
8.4.7.1	Using Regular Expressions in a Whitelist	158
8.4.7.2	Specifying Literal Dot Separators	158
8.4.7.3	Specifying Domains.....	158
8.4.7.4	Specifying Host Names	159
8.4.8	Configure MashZone NextGen for SSL and Digital Certificates.....	159
8.4.8.1	The Certificate Store and Certificates.....	160
8.4.8.2	Key Certificate Pairs	160
8.4.8.3	Trusted Peer Certificates.....	160
8.4.8.4	The Certificate Store	160
8.4.8.5	Configure Mutual SSL Between Users and MashZone NextGen	161
8.4.8.6	One-Way SSL to MashZone NextGen	161
8.4.8.7	One-Way SSL to Information Sources.....	161
8.4.8.8	Configure HTTPS and Certificate Stores in the Application Server.....	162
8.4.8.9	Update SSL Configuration for Java	163
8.4.9	MashZone NextGen Logging	164
8.4.9.1	Configure Logging for the MashZone NextGen Server	164
8.4.9.2	Audit logging for dashboards, data feeds, aliases, and permissions	165
8.4.10	MashZone NextGen Notifications.....	166
8.4.10.1	Configuring a Mail Server for MashZone NextGen	166
8.4.10.2	Update the User Email Attribute from LDAP	167
8.4.11	BigMemory for Caching, Connections and In-Memory Stores	167
8.4.11.1	Caching for the MashZone NextGen Server.....	168
8.4.11.1.1	Artifact Caching	169
8.4.11.1.2	Response Caching	169
8.4.11.1.3	Distributed Caching for MashZone NextGen Clusters.....	169
8.4.11.1.4	Configure BigMemory Servers for MashZone NextGen Caching	170
8.4.11.2	Working with BigMemory Stores for RAQL	171
8.4.11.2.1	Declared In-Memory Stores.....	171
8.4.11.2.1.1	Declare a new In-Memory Store.....	172
8.4.11.2.1.2	Modify a Declared In-Memory Store.....	174
8.4.11.2.1.3	View Details for Declared In-Memory Stores	174
8.4.11.3	Dynamic In-Memory Stores.....	174
8.4.11.3.1	Manage Dynamic BigMemory Stores for MashZone NextGen Analytics	175
8.4.11.3.2	Add an External Dynamic In-Memory Store Connection.....	175
8.4.11.3.3	Delete External Dynamic In-Memory Store Connections	176
8.4.12	Manage Terracotta DB connections	177
8.4.12.1	Register Terracotta DB connections	177
8.4.12.2	Edit Terracotta DB connections	178
8.4.12.3	Test Terracotta DB connections.....	179
8.4.12.4	Share Terracotta DB connections	179
8.4.12.5	Delete Terracotta DB connections	180

8.4.13	Manage data sources and drivers	180
8.4.13.1	Add a data source	180
8.4.13.2	Edit, test or remove data sources.....	182
8.4.13.3	Share data sources	183
8.4.13.4	Add or manage JDBC drivers	183
8.4.13.5	Migrate JDBC connections	184
8.4.13.5.1	Migrate JDBC configuration of Presto to MashZone NextGen ...	184
8.4.13.5.2	Migrate JDBC connections of Presto to MashZone NextGen	184
8.4.13.5.3	Migrate JDBC configuration of MashZone NextGen 9.10	185
8.4.13.5.4	Migrate JDBC connections of MashZone legacy to MashZone NextGen	186
8.4.14	Manage map files.....	186
8.4.14.1	Manage geoJSON files	186
8.4.14.2	Manage tile server configuration files.....	187
8.4.15	Tune memory/caching for MashZone NextGen	187
8.4.15.1	Tune MashZone Memory and Cache Configuration Manually	188
8.4.15.2	Update Cache Memory Settings	188
8.4.15.3	Update MashZone ThreadSize Properties.....	188
8.4.16	Manage Materialized Feeds.....	188
8.4.16.1	Configure Materialized Feeds	189
8.4.16.2	Edit Materialized Feeds	189
8.4.16.3	Delete Materialized Feeds.....	190
8.4.16.4	Share Materialized Feeds	190
8.4.16.5	Materialize data feed results in Terracotta DB	191
8.4.16.5.1	Use the Admin Console	191
8.4.16.5.2	Use the API Console	192
8.4.16.5.3	Use the server administration tool.....	192
8.4.16.5.4	Use server web services.....	193
8.5	MashZone NextGen Server Administration	195
8.5.1	Manage Files for MashZone NextGen Features or Artifacts.....	195
8.5.1.1	Add External Resources as MashZone NextGen Files.....	195
8.5.1.2	Find MashZone NextGen Files	196
8.5.1.3	Update or Delete MashZone NextGen Files	196
8.5.2	Manage resource directories.....	196
8.5.2.1	Create resource directory.....	197
8.5.2.2	Change resource directory	197
8.5.2.3	Delete resource directory	197
8.5.2.4	Share resource directory	198
8.5.3	Manage URL aliases	198
8.5.3.1	Create URL alias	198
8.5.3.2	Change URL alias	199
8.5.3.3	Delete URL alias.....	199
8.5.3.4	Share URL alias	200
8.5.4	Deploying MashZone NextGen Instances, Clusters or Artifacts.....	200
8.5.4.1	Deploying the Core Widgets	200
8.5.4.2	Deploying MashZone NextGen Artifacts and Other Metadata	201
8.5.4.3	Export users, user metadata and groups.....	204
8.5.4.4	Export dashboards	205
8.5.4.5	Export data feeds.....	205
8.5.4.6	Import users, user metadata and groups	206
8.5.4.7	Import dashboards.....	207
8.5.4.8	Import data feeds	208
8.5.4.9	Deploying multiple MashZone NextGen servers in one host	209

8.5.5	Clustering MashZone NextGen Servers	209
8.5.5.1	Setting Up a New Cluster	209
8.5.5.2	Adding New Members to an Existing Cluster	211
8.5.6	Sharing the MashZone NextGen Repository in Clustered Environments	212
8.5.6.1	Create and Share a New MashZone NextGen Repository	212
8.5.6.2	Share an Existing MashZone NextGen Repository	213
8.5.7	Setting Up an External MashZone NextGen Configuration Folder	213
8.5.7.1	MashZone NextGen File-Based Configuration and Extensions	215
8.5.7.2	MashZone NextGen Configuration Files That Can Be External	215
8.5.7.3	MashZone NextGen Configuration Files That Must Be Internal	216
8.5.7.4	MashZone NextGen Extensions	217
8.5.8	MashZone NextGen dashboards in a clustered scenario	217
8.5.8.1	Preliminary	217
8.5.8.2	Configuration	218
8.5.8.2.1	MashZone NextGen Event Service	218
8.5.8.3	rtbs.properties	218
8.5.8.4	zookeeper.properties	219
8.5.8.5	server.properties	220
8.5.8.6	MashZone NextGen nodes	220
8.5.8.6.1	Customizing dashboards	220
8.5.8.7	Custom styles	220
8.5.8.8	Custom widgets	221
8.5.8.9	Using JDBC drivers	221
8.5.8.10	Local file resources	221
8.6	Event Service Configuration and Administration	221
8.6.1	About the Event Service and Event Data	221
8.6.2	Use Events as Information Sources	222
8.6.3	Manage EDA Event Sources	222
8.6.3.1	Create EDA Event Sources	223
8.6.3.2	Edit EDA Event Sources	227
8.6.3.2.1	Simple and Hybrid Throttling Strategies	230
8.6.3.3	Duplicate EDA Event Sources	231
8.6.3.4	Delete EDA Event Sources	231
8.6.3.5	Share EDA Event Sources	232
8.6.4	Manage Apama Event Sources	232
8.6.4.1	Create Apama Event Sources	233
8.6.4.2	Duplicate Apama Event Sources	238
8.6.4.3	Delete Apama Event Sources	238
8.6.4.4	Edit Apama Event Sources	238
8.6.4.5	Share Apama Event Sources	243
8.6.5	Manage DES Event Sources	244
8.6.5.1	Create DES Event Source	244
8.6.5.2	Edit DES Event Sources	249
8.6.5.2.1	Simple and Hybrid Throttling Strategies	252
8.6.5.3	Duplicate DES Event Sources	253
8.6.5.4	Delete DES Event Sources	253
8.6.5.5	Share DES Event Sources	253
8.6.5.6	Activate DES in MashZone NextGen	254
8.6.6	Start or Stop an Event Source	254
8.6.7	Restart all Event Sources	255
8.6.8	Manage Apama Instances	255

8.6.8.1	Create Apama Instances	255
8.6.8.2	Edit Apama Instances	256
8.6.8.3	Delete Apama Instances	256
8.6.9	Manage Apama Event Targets	257
8.6.9.1	Create Apama Event Targets	257
8.6.9.2	Edit Apama Event Targets	258
8.6.9.3	Delete Apama Event Targets	259
8.6.9.4	Share Apama Event Target.....	259
8.7	Process Performance Manager Integration	260
8.7.1	Manage PPM Connections.....	260
8.7.2	Create PPM Connections	261
8.7.3	Edit PPM Connections	262
8.7.4	Delete PPM Connections.....	263
8.7.5	Share PPM connections	263
8.8	webMethods Business Console Integration	264
8.8.1	Example	264
8.8.2	Authentication	264
8.8.3	Example URL.....	265
8.8.4	Configuration	265
8.8.5	Outbound API.....	265
8.8.6	Inbound API	266
8.9	MashZone NextGen Repositories.....	266
8.9.1	Maintenance Suggestions.....	266
8.9.2	Tuning the MashZone NextGen Repository Connection Pool	267
8.9.2.1	Connection Pool Size Properties	267
8.9.2.2	Idle Pool Connection Properties.....	267
8.9.3	Synchronize the MashZone NextGen Repository and MashZone NextGen Server Time Zones.....	268
9	Appendix.....	269
9.1	Widgets.....	269
9.1.1	Action button	269
9.1.2	Bar chart	269
9.1.3	Bubble chart	271
9.1.4	Circular gauge chart.....	272
9.1.5	Column chart	273
9.1.6	Date filter	275
9.1.7	Drop-down box.....	276
9.1.8	Grid	277
9.1.9	Horizontal and vertical gauge chart	279
9.1.10	Image	280
9.1.11	Input field.....	281
9.1.12	Label.....	282
9.1.13	Line chart	283
9.1.14	List	284
9.1.15	Map with markers	286
9.1.16	Pie chart	287
9.1.17	Rich text area.....	288
9.1.18	Slider	290
9.1.19	Jump to PPM client.....	292
9.1.20	Traffic lights.....	294
9.1.21	Vector map	295
9.2	Operators	296
9.2.1	Data source operators	296

9.2.1.1	ARIS table.....	296
9.2.1.2	BigMemory.....	297
9.2.1.3	CSV.....	299
9.2.1.4	Data feed.....	302
9.2.1.5	Events.....	302
9.2.1.6	JDBC.....	303
9.2.1.7	JSON.....	304
9.2.1.8	Excel.....	306
9.2.1.9	PPM.....	309
9.2.1.9.1	Add filter.....	310
9.2.1.10	Terracotta DB.....	311
9.2.1.11	XML.....	313
9.2.1.12	URL parameter syntax.....	315
9.2.1.13	RAQL and SQL statement parameters.....	316
9.2.1.13.1	JDBC parameters.....	316
9.2.1.13.2	RAQL parameters.....	316
9.2.2	Data transformation operators.....	317
9.2.2.1	Aggregate.....	317
9.2.2.1.1	Parameters.....	317
9.2.2.1.2	Aggregation type.....	318
9.2.2.1.3	Examples.....	319
9.2.2.2	Arithmetic.....	319
9.2.2.2.1	Parameters.....	320
9.2.2.2.2	Example.....	321
9.2.2.3	Average.....	321
9.2.2.3.1	Parameters.....	321
9.2.2.4	Change data type.....	322
9.2.2.4.1	Parameters.....	323
9.2.2.4.2	Examples.....	327
9.2.2.5	Change data type - single value.....	327
9.2.2.5.1	Parameters.....	329
9.2.2.5.2	Examples.....	332
9.2.2.6	Column to value.....	333
9.2.2.6.1	Parameters.....	333
9.2.2.7	Combine data feeds.....	333
9.2.2.7.1	Parameters.....	333
9.2.2.8	Concatenate data feeds.....	334
9.2.2.8.1	Parameters.....	334
9.2.2.9	Concatenate texts.....	335
9.2.2.9.1	Parameters.....	335
9.2.2.10	Conditional replace.....	335
9.2.2.10.1	Parameters.....	336
9.2.2.10.2	Comparison operators.....	337
9.2.2.11	Convert text.....	337
9.2.2.11.1	Parameters.....	337
9.2.2.12	Copy data feeds.....	338
9.2.2.12.1	Parameters.....	338
9.2.2.13	Copy single value.....	339
9.2.2.13.1	Parameters.....	339

9.2.2.14	Delete column	339
9.2.2.14.1	Parameters.....	339
9.2.2.15	Duplicate column	339
9.2.2.15.1	Parameter	340
9.2.2.16	Extract text	340
9.2.2.16.1	Parameters.....	340
9.2.2.17	Filter by date.....	341
9.2.2.17.1	Parameters.....	341
9.2.2.18	Filter rows.....	342
9.2.2.18.1	Parameters.....	343
9.2.2.19	Find text index	344
9.2.2.19.1	Parameters.....	344
9.2.2.20	Goal accomplishment	345
9.2.2.20.1	Parameters.....	345
9.2.2.21	Insert column	346
9.2.2.21.1	Parameter	346
9.2.2.22	Merge single texts.....	347
9.2.2.22.1	Parameters.....	347
9.2.2.23	Move date	348
9.2.2.23.1	Parameter	348
9.2.2.23.2	Example.....	349
9.2.2.24	RAQL Inline.....	350
9.2.2.24.1	Parameters.....	350
9.2.2.24.2	Example.....	350
9.2.2.25	Rename column	350
9.2.2.25.1	Parameter	350
9.2.2.26	Replace text	351
9.2.2.26.1	Parameters.....	351
9.2.2.27	Round up/down	352
9.2.2.27.1	Parameter	352
9.2.2.28	Round up/down date	353
9.2.2.28.1	Parameter	353
9.2.2.28.2	Examples	354
9.2.2.29	Runtime info	355
9.2.2.29.1	Parameters.....	355
9.2.2.30	Value to column.....	356
9.2.2.30.1	Parameters.....	356
9.2.3	User input operators	357
9.2.3.1	Date user input.....	357
9.2.3.1.1	Parameters.....	357
9.2.3.2	Date user input (List)	358
9.2.3.2.1	Parameters.....	358
9.2.3.3	Number user input	359
9.2.3.3.1	Parameters.....	359
9.2.3.4	Number user input (List).....	359
9.2.3.4.1	Parameters.....	360

9.2.3.5	Text user input	360
9.2.3.5.1	Parameters.....	360
9.2.3.6	Text user input (List)	361
9.2.3.6.1	Parameters.....	361
9.3	Analytics with RAQL	362
9.3.1	About the Real-Time Analytical Query Language	362
9.3.1.1	Scope of RAQL.....	362
9.3.1.2	RAQL in MashZone NextGen	362
9.3.2	Working with RAQL queries	363
9.3.2.1	Query Syntax of RAQL	363
9.3.2.2	Data types in RAQL	364
9.3.2.3	Creating dynamic RAQL queries	365
9.3.3	Language reference	365
9.3.3.1	Reserved keywords	365
9.3.3.2	Syntax Reference.....	367
9.3.3.2.1	ExplainOrSelectStatement	367
9.3.3.2.2	CharacterConstant	367
9.3.3.2.3	SelectStatement	368
9.3.3.2.4	SelectWithOrder.....	368
9.3.3.2.5	QueryExpressionBody	368
9.3.3.2.6	QueryTerm.....	369
9.3.3.2.7	QueryPrimary	369
9.3.3.2.8	SelectWithoutOrder	369
9.3.3.2.9	SelectList	370
9.3.3.2.10	SelectItem	370
9.3.3.2.11	SelectStar	370
9.3.3.2.12	ObjectName	371
9.3.3.2.13	SqlSimpleExpression	371
9.3.3.2.14	SqlMultiplicativeExpression.....	372
9.3.3.2.15	SqlExponentExpression	372
9.3.3.2.16	SqlUnaryExpression.....	372
9.3.3.2.17	SqlPrimaryExpression	373
9.3.3.2.18	SqlExpression	373
9.3.3.2.19	SqlAndExpression.....	374
9.3.3.2.20	SqlUnaryLogicalExpression	374
9.3.3.2.21	SqlRelationalExpression	375
9.3.3.2.22	SqlExpressionList	375
9.3.3.2.23	SqlRelationalOperatorExpression	376
9.3.3.2.24	Relop.....	376
9.3.3.2.25	SqlInClause	376
9.3.3.2.26	SqlBetweenClause.....	377
9.3.3.2.27	SqlLikeClause	377
9.3.3.2.28	IsNullClause	377
9.3.3.2.29	Constant	378
9.3.3.2.30	FunctionCall	379
9.3.3.2.31	CastableDataTypeDeclaration	381
9.3.3.2.32	UserDefinedType.....	382
9.3.3.2.33	SimpleCaseCall	382
9.3.3.2.34	BooleanCaseCall	382
9.3.3.2.35	AggregateFunctionCall	384
9.3.3.2.36	SourceField	385
9.3.3.2.37	GenericFunctionCall.....	385
9.3.3.2.38	ConstantList	386
9.3.3.2.39	WindowSpecification.....	386

9.3.3.2.40	OrderByClause.....	386
9.3.3.2.41	WindowFrameClause	387
9.3.3.2.42	WindowFrameStart	387
9.3.3.2.43	FrameBound	387
9.3.3.2.44	WindowFrameBound	388
9.3.3.2.45	GroupingOperation	388
9.3.3.2.46	WindowFunctionCall	389
9.3.3.2.47	LongConstant	389
9.3.3.2.48	IntegerConstant.....	390
9.3.3.2.49	ArrayElementReference.....	390
9.3.3.2.50	ArrayValueConstructor	390
9.3.3.2.51	Alias	391
9.3.3.2.52	FromClause	391
9.3.3.2.53	FromItem.....	391
9.3.3.2.54	SourceReferenceOrSubquery	391
9.3.3.2.55	SourceReference	392
9.3.3.2.56	JoinedTable	392
9.3.3.2.57	SourceFieldList	393
9.3.3.2.58	WhereClause	393
9.3.3.2.59	GroupByClause	393
9.3.3.2.60	GroupingElementList	394
9.3.3.2.61	GroupingElement	394
9.3.3.2.62	EmptyGroupingSet	394
9.3.3.2.63	OrdinaryGroupingSet.....	395
9.3.3.2.64	GroupingExpressionReferenceList.....	395
9.3.3.2.65	GroupingExpressionReference.....	395
9.3.3.2.66	RollupList.....	396
9.3.3.2.67	OrdinaryGroupingSetList	396
9.3.3.2.68	CubeList.....	396
9.3.3.2.69	GroupingSetsSpecification	396
9.3.3.2.70	GroupingSetList	397
9.3.3.2.71	GroupingSet	397
9.3.3.2.72	HavingClause.....	397
9.3.3.2.73	LimitClause	398
9.3.3.2.74	S_NUMBER.....	398
9.3.3.2.75	FLOAT.....	398
9.3.3.2.76	INTEGER.....	399
9.3.3.2.77	DIGIT	399
9.3.3.2.78	S_PARAMETER.....	399
9.3.3.2.79	S_IDENTIFIER	400
9.3.3.2.80	NAMESTARTCHAR	401
9.3.3.2.81	NAMECHAR.....	402
9.3.3.2.82	S_QUOTED_IDENTIFIER.....	402
9.3.3.2.83	S_CHAR_LITERAL.....	403
9.3.3.2.84	IgnorableWhitespace	403
9.3.3.2.85	LINE_COMMENT	403
9.3.3.2.86	MULTI_LINE_COMMENT	404
9.3.4	RAQL extensions.....	404
9.3.4.1	RAQL function extensions.....	404
9.3.4.1.1	Plain functions	405
9.3.4.1.2	Analytic functions.....	409
9.3.4.1.3	Window functions over RANGES	413
9.3.4.1.4	Migrate RAQL queries from Version 3.7 to 3.8.....	416
9.3.4.1.4.1	Update plain functions for 3.8	417
9.3.4.1.4.2	Update analytic functions for 3.8.....	418
9.3.4.2	Creating new user-defined functions.....	419

9.3.4.2.1	Set Up Your Development Environment.....	420
9.3.4.2.2	Write Plain Functions for RAQL.....	421
9.3.4.2.3	Configure, Compile, Deploy and Test User-Defined Functions	422
9.3.4.2.3.1	UDF Library Configuration	424
9.3.4.2.3.2	External UDF Library Deployment Folder	425
9.3.4.2.4	Write Window Analytic Functions for RAQL	425
9.3.4.2.4.1	Construct and Initialize the Window Analytic Function Class	426
9.3.4.2.4.2	Implement the call Method for the Window Analytic Function.....	427
9.3.4.2.4.3	Complete WindowSum Example	429
9.3.4.2.4.4	Work with Specific Records in Window Calculations.....	430
9.3.4.2.5	Write Aggregate Analytic Functions for RAQL	432
9.3.4.2.5.1	My Average Aggregate Example	433
9.3.4.2.5.2	Kurtosis Using a Third Party Library	435
9.3.4.2.6	Use a Factory for Function Overloading	437
9.4	MashZone NextGen Platform API Console	438
9.4.1	Open the MashZone NextGen API Console	438
9.4.2	MashZone NextGen Platform Services.....	438
9.4.3	Input Parameters for JUMP Requests	439
9.4.3.1	Parameter Values.....	440
9.5	Root Cause Miner.....	441
9.5.1	Symptoms and root causes	441
9.5.2	Limitations	442
9.5.3	Score	442
9.5.4	Scaled system	443
9.5.5	Permissions.....	443
9.5.6	How to use the Root Cause Miner widget	443
10	Additional Information and Support	444
10.1	Samples, Help and Other Documentation	444
10.2	Version and License Information	444
11	Legal information.....	445
11.1	Documentation scope	445
11.2	Data protection	445

1 Preface

The Software AG MashZone NextGen User & Developer Guide provides an introduction and instructions for MashZone NextGen business users and developers. This guide discusses topics for MashZone NextGen version 10.3.

TOPICS

Introducing Software AG MashZone NextGen (Seite 2)

MashZone NextGen welcome page (page 3)

Open a dashboard in MashZone NextGen (Seite 4)

Use the dashboard editor (page 5)

Use the data feed editor (page 81)

Administration (Seite 95)

Appendix (Seite 269)

Additional Information and Support (page 444)

2 Introducing Software AG MashZone NextGen

Software AG MashZone NextGen provides self-service analytics for business users.

MashZone NextGen is a browser-based application enabling you to analyze and visualize any data from various, independently distributed data sources. The data sources that were combined using data feeds (page 93) are represented graphically and analyzed on dashboards (page 93).

IN-MEMORY ANALYTICS

- MashZone NextGen leverages Terracotta BigMemory to offer an in-memory analytics language, Event Service, for analyzing small and large amounts of data.
- Visualize data kept in Terracotta cache on MashZone NextGen dashboards
- Use Terracotta cache as an in-memory store of MashZone NextGen where data can be held in-memory on locally on-heap, locally off-heap or remote.

VISUALIZATION & DATA MASHING

- Create user-friendly dashboards using the intuitive user interface.
- Rely on comprehensive integration support for enterprise standards, including JDBC® support, CSV files, XML files, BigMemory, Apama and ARIS tables.

COMPREHENSIVE LIBRARY OF WIDGETS

All MashZone NextGen dashboards are based on HTML5 standards. Use these enhanced widgets to create dashboards, for example, bar chart, grids, traffic light.

3 MashZone NextGen welcome page

The MashZone NextGen welcome page gives you a quick access to your dashboards and data feeds. And it allows you to manage your user data and system settings in the MashZone NextGen Administration.

To display the MashZone NextGen welcome page, click the MashZone NextGen logo in the program bar.

On the welcome page, you have several possibilities to open and to create your dashboards and data feeds:

- The **DASHBOARDS** and **DATA FEEDS** menu in the MashZone NextGen program bar allows you to open and to create dashboards and data feeds.
- The **Recent Dashboards** area shows the user's most recently edited dashboards.
- The **Search** box allows you to search all dashboards and data feeds available in MashZone NextGen.

More functions are available in the user menu of the program bar. To open the user menu, click the  user icon in the program bar.

- **My profile** enables you to edit your user data.
- **Admin Console** opens the MashZone NextGen Administration.
- **Help?** displays the MashZone NextGen online help.
- **Logout** logs you out from MashZone NextGen.
- **API Console** opens the MashZone NextGen API console.

4 Open a dashboard in MashZone NextGen

MashZone NextGen provides two modes for opening your dashboards. In view mode, you can open your dashboards with current data for interactive use. In edit mode, you can open available dashboards to edit their layout or data assignments, for example,.

You need the appropriate privileges to open a dashboard in view or edit mode.

Procedure

1. Click **Software AG MashZone NextGen** in the program bar to open the MashZone NextGen welcome page.
2. Click a dashboard in the **Recent Dashboard** list.
The list contains all your dashboards that you recently viewed.
3. To find a dashboard, enter a dashboard name or a part of the name in the **Search** box and click **Search**. In the search result list, click a dashboard to open it.
The dashboard opens in view mode.
4. Click **Dashboards > Open dashboard** in the MashZone NextGen program bar.
 - a. Select an available dashboard. You can also search a dashboard using your keyboard.
 - b. Click **OK**.

The dashboard opens in view mode.

You can switch from view mode to edit mode and vice versa at any time, provided you have the appropriate privileges. To switch from view mode to edit mode, click the  **Edit dashboard** icon in the dashboard main menu. To switch from edit mode to view mode, click the  **View dashboard** icon in the dashboard main menu.

5 Use the dashboard editor

The dashboard editor as a graphical user interface is an easy-to-use tool to create, manage, and view your dashboards. The dashboard editor offers edit mode and view mode. In edit mode, you can create, edit and manage your dashboards. The view mode allows you to view and use your dashboards interactively.

- To use the edit mode, the **Dashboard administrator** function privilege is required.
- To display a dashboard in view mode, the **View** permission is required. See Manage dashboard permissions (page 16) for details.

5.1 Use dashboards in view mode

You can display and use your dashboards interactively in view mode.

To use a dashboard in view mode, you need the **View** permission for the dashboard. See Manage dashboard permissions (page 16) for details.

5.1.1 Display a dashboard in view mode

You can switch from edit mode to view mode in the dashboard editor. Your currently opened dashboards are displayed with real data in view mode.

Procedure

1. Create (page 11) or open (page 10) a dashboard in edit mode in the dashboard editor.
2. Click  **View dashboard** in the dashboard editor main menu.

The dashboard is displayed with real data in view mode in the dashboard editor.

To switch from view mode to edit mode, click  **Edit dashboard** in the main menu.

5.1.2 Use interactive filters in dashboards

You can use interactive filters for charts and tables if the relevant filter conditions have been defined for these widgets. You can set a filter by clicking a data element of a widget, for example, a data point in a line chart. The selected data element of this widget takes effect as a filter for all associated widgets.

MULTIPLE SELECTION

Some widgets allow you to select multiple values at the same time to filter other widgets. You can select, for example, multiple rows in a table or multiple data points in a chart. Multiple selection is available if the corresponding **Multiple selection** option is enabled for a widget. The widgets providing the **Multiple selection** option are Column chart (page 273), Bar chart (page 269), Pie

chart (page 287), Grid (page 277), and List (page 284). The widget to be filtered must also support multiple selection. For details, see Define filters for widgets (Seite 60).

Procedure

1. Display a dashboard in dashboard view mode. (Seite 5)
2. Click a data element of a widget, for example, a coordinate of a line chart or a cell in a table. The selected data element is applied as a filter to all associated widgets.
3. Click the  **Menu** icon > **Clear selection** to clear your filter settings. In charts, you can also click in the background of a widget to cancel the selection of your filter. The menu is available if the **Show menu** option is enabled for the widget.

The widgets are displayed with the filters selected.

You can define filters for specific widgets. See Define filters for widgets (Seite 60).

5.1.3 Use the filter panel of Process Mining context-based widgets

You can display and configure the filter panel in dashboard edit and view mode.

Using the filter panel, you can filter widgets of a dashboard whose data are based on a Process Mining context. The filter panel replaces the standard filter method used for non-context-based widgets (Seite 5). The filter provides all text-based and filterable numeric data columns that are not yet assigned to any other context-based widget.

- Text-based filter criteria are displayed as list.
- Numeric filter criteria are provided as slider in the filter panel. The basic settings are automatically preconfigured, these are range, initial selection, step width, and numeric format. You can manually adjust the minimum and maximum values of the slider widget. See **Slider (page 290)** widget for details.

Prerequisites

The filter panel is enabled (page 24).

Procedure

1. Open a dashboard in dashboard view mode.
2. Click the << icon to show the filter panel on the right-hand side of the dashboard. The panel is hidden by default and no filter columns are specified.
3. To hide the filter panel, click the >> icon.
4. If no filter is specified, add filter criteria to the filter panel.
 - a. Click the  **Add filter** icon to add an initial filter.
 - b. Click the  **Filter settings** icon to display filter options.
 - c. Enter a term in the  **Search** box to filter the column list.

The list includes all data columns of text type that have not been assigned as filter columns to the Process Mining context-based widgets. See Assign data columns to Process Mining context-based widgets (page 21).

- d. Enable the **Selected** option to display only the data columns that have already been selected.
- e. Select the data columns you want to use as filter criteria.
- f. Click **OK**.

Your settings are applied. An appropriate filter criterion is added to the filter panel for each selected data column.

5. Specify further filter criteria to be used in the panel.
 - a. Click **Edit filter panel**.
 - b. Click the  **Filter settings** icon to display filter options.
 - c. Enter a term in the  **Search** box to filter the values of a filter criterion.
 - d. Enable the **Selected** option to display only the filter values that have already been selected.
 - e. Select the data columns you want to use as filter criteria.
 - f. You can use the filter criteria configured for context-based widgets (page 22) in the filter panel. As long as such a filter criterion is enabled in the filter panel, the filter criterion is not longer available in the widget. These filter criteria are hidden by default.
 1. Enable the option **Show configured context filters from dashboard widgets**. The filter criteria are displayed and marked in bold.
 2. Select the context filter criteria that you want to use in the filter panel.
 - g. Click **OK**.

Your settings are applied. An appropriate filter criterion is added to the filter panel for each selected data column.
6. Click the name of a text-base filter criterion to sort the values of the criterion.
7. Select the individual values of the text-base criteria by which you want to filter the widgets.
8. Adjust the slider using the mouse pointer to select the numeric filter values. You can adjust the minimum and maximum values and you can move the entire selected data range with the mouse pointer.

The filter panel is configured and the filter values for the widgets are specified. The values of all widgets are filtered according to your filter settings.

5.1.4 Refresh data of widgets

You can manually refresh the data currently displayed in a widget.

The data is extracted from the cache or recalculated if the refresh rate of the data source has expired. The default refresh rate value is 12 h. You can set the **Refresh rate** in the Assign data (Seite 31) dialog.

The use of the manual **Refresh** option is independent of the **Auto Refresh** function.

Procedure

1. Display a dashboard in dashboard view mode. (Seite 5)
2. Click a widget.
3. Click the  **Menu** icon > **Refresh**.

The widgets are displayed with the refreshed data.

Most of the widgets provide the **Auto refresh** option. Use this option to enable automatic data retrieval for a widget. The source data is reimported and recalculated automatically based on the refresh rate set. The **Auto refresh** option is available on the **Config** tab in the widget properties.

To display the properties dialog in dashboard edit mode () , click an inserted widget.

5.1.5 Pause automatic data refreshing

You can pause the automatic refreshing of widget data.

The data is automatically extracted from the cache or recalculated if the refresh rate of the data source has expired. The default refresh rate value is 12 h. You can set the **Refresh rate** in the Assign data (Seite 31) dialog.

Procedure

1. Display a dashboard in dashboard view mode. (Seite 5)
2. Click a widget.
3. Click the  **Menu** icon > **Pause** to pause the automatic data update.
4. Click the  **Menu** icon > **Resume** to restart the automatic data update.

The automatic refreshing of widget data is paused.

Most of the widgets provide the **Auto refresh** option. Use this option to enable automatic data retrieval for a widget. The source data is reimported and recalculated automatically based on the refresh rate set. The **Auto refresh** option is available on the **Config** tab in the widget properties.

To display the properties dialog in dashboard edit mode () , click an inserted widget.

5.1.6 Change column width and sort order

In dashboard view mode, you can change the initial table column width and sort order of a **Grid** widget.

Procedure

1. Open a dashboard in view mode. (Seite 5)
2. Click a table column header of an inserted **Grid** widget.

The first click changes the sort order to 'ascending', the second click to 'descending', and the third click to 'unsorted' if this was the initial state.

You can sort several columns simultaneously. To select more than one column, press the **Shift** key and click the required column headings. The rows are sorted by the first column header.

3. To change the width of a column, drag the boundary on the right side of the column header accordingly using the mouse.

Your settings are applied.

5.1.7 Use lists for multiple selection

You can use the List widget for the multiple selection of values.

In particular, a dashboard can provide lists for selecting individual or multiple values. The List widget (page 284) provides the multiple selection option by default. The values selected are mainly used to filter values of other widgets (Seite 5).

A list can contain one or two columns. The main column on the left can be used to sort the list, for example, it can contain product names. The second column provides only additional values, such as the product prices. A list also provides check boxes to select multiple values if the **Multiple selection** option is enabled for the widget.

Procedure

1. Open a dashboard in view mode. (Seite 5)
2. Click a row in a **List** widget to select the contained value.
3. Click a row, press the Shift key, and click another row to select multiple values at the same time.
4. Click the check box in the column header to select all values in the list. Click the check box again to clear all list values.
5. Click the name in the header of the main column to sort the list.

The first click changes the sort order to ascending, the second click to descending, and the third click to unsorted, if this was the initial state.

6. You can filter the list values.
 - a. Click the  **Search** icon in the list header.
 - b. Enter your filter term in the **Search** box.
 - c. Click **Selected** to show only the selected values in the filtered list.

Your settings are applied.

Multiple selection is also available in the widgets Column chart (page 273), Bar chart (page 269), Pie chart (page 287), and Grid (page 277).

Assign data columns to lists

5.1.8 Save widget data as a CSV file

You can save the current data visualized in a widget as a CSV file.

- The data are saved in the default data format or in the format set in data assignment (Seite 31) (Assign data 2/2 dialog).
- The aggregation and the value rounding are taken from the column configuration set in data assignment (Seite 31) on page 78 (Assign data 2/2 dialog).
- The sorting is taken from the data feed result based on the feed definition.
- The default separator is ,.
- The default masking is ".

Procedure

1. Open a dashboard in dashboard view mode. (Seite 5)
2. Click a widget.
3. Click the  **Menu** icon.

The menu is available if the **Widget menu** option is enabled for the widget.

4. Click **Save as CSV**.
5. Make your settings.

Depending on which Web browser you are using, you can select an application that you want to open to view the data, or you can save the data directly as a CSV file.

The widget data is saved as a CSV file.

Note that CSV files can pose a security risk when they are opened in MS Excel. Certain characters in the CSV file can be used for unwanted code execution.

Enclose all values beginning with =, +, -, or @ in single quotation mark before exporting as CSV file. For details, refer to

<https://www.contextis.com//resources/blog/comma-separated-vulnerabilities/>.

5.2 Manage dashboards

In the dashboard editor, you can manage your dashboards in edit mode.

5.2.1 Open a dashboard in the dashboard editor

In the dashboard editor, you can open an existing dashboard in edit mode.

Procedure

1. Click **Manage** > **Open** in the dashboard editor main menu.
2. Select a dashboard in the list of available dashboards.
You can also search for a dashboard using your keyboard.

3. Click **OK**.

The selected dashboard opens in the dashboard editor.

5.2.2 Create a dashboard in the dashboard editor

In the dashboard editor, you can create and configure your dashboards.

Procedure

1. Click **Manage** > **Create dashboard** in the dashboard editor main menu.
A new dashboard opens.
2. Configure your dashboard.
3. Click **Manage** > **Save** in the dashboard main menu. The **Properties** dialog of the dashboard opens.
4. Enter a dashboard name in the **Name** box.
5. Optionally, enter a dashboard description in the **Description** box.
6. Optionally, enter comma-separated search tags in the **Tags** box. The search tags help you to find your dashboard when you are using the search function.
7. Click **OK**.

Your changes are applied.

You can edit the dashboard properties (Seite 16) later as needed.

To display the dashboard in view mode, click  **View dashboard** in the main menu.

5.2.3 Use smart dashboard edit mode

The edit mode of the dashboard editor is an easy-to-use tool to create, manage and edit your dashboards.

The edit mode provides two different dashboard workspaces, that is, a smart dashboard and a fixed-grid dashboard.

Unlike a fixed-grid dashboard, the size of a smart dashboard is automatically adapted to the screen resolutions in dashboard view mode. Depending on the space available, widgets are stretched or compressed. In case there is not enough space available, widgets are automatically re-positioned. See Use dashboards in view mode (page 5) for details.

The smart dashboard workspace is displayed by default when you create a new dashboard (page 11) or when you open an existing smart dashboard.

The workspace is divided in three rows and 12 columns in which you can place your widgets (page 30).

5.2.3.1 Insert new row

You can insert a new row in the dashboard.

Procedure

1. Click a row in the smart dashboard workspace.
2. Click **Insert row** in the properties dialog of the row.

A new line is inserted below the row selected.

5.2.3.2 Delete a row

You can delete a row from the dashboard.

Procedure

1. Click a row in the smart dashboard workspace.
2. Click **Delete** in the properties dialog of the row.

The row is deleted.

5.2.3.3 Resize a row

You can change the height of a row.

Procedure

1. Move the cursor to the lower boundary of a row.
The cursor will change to a resize symbol.
2. Click and drag the row boundary using your mouse.
The height of all widgets located in the row is automatically resized.

The row height is resized.

5.2.3.4 Lock row height

You can lock the row height set in the dashboard workspace.

By default, the row height is adapted dynamically to the screen resolution of the dashboard view mode and therefore the size of the widgets contained in the dashboard is adapted as well. If you lock the row height, the widget size is also fixed.

Procedure

1. Click a row in the workspace.
2. Enable the **Lock height** option in the properties dialog of the row.

The row height is fixed.

5.2.3.5 Change a row order

You can change the row order on the smart dashboard.

Procedure

1. Click a row in the workspace.
2. Move the row selected up or down using drag and drop

The row order is changed.

5.2.3.6 Place a widget

You can place a widget in any empty field in the dashboard workspace.

Procedure

1. Click a widget in the dashboard workspace.
2. Move the widget selected using drag and drop and place it in an empty field on the dashboard.

You cannot place more than one widget in one field.

The widget is placed in an empty field.

5.2.3.7 Resize widgets

You can scale the size of widgets up or down.

Procedure

1. Click a widget on the dashboard. The widget is displayed with a frame.
2. Resize the widget width by dragging the anchor point of the frame with your mouse pointer. A widget width can be resized across multiple empty fields.
3. Resize the widget height by resizing the height of the row containing the widget. To resize the row height, drag the upper or lower row border with your mouse pointer.

The height of all widgets inserted into the same row is resized automatically.

The selected widgets are resized.

5.2.3.8 Avoid line break in dashboard view mode

You can avoid the automatic line break of the widgets if there is not enough space available in dashboard view mode.

Procedure

1. Click a row in the workspace.
2. Enable the **Do not break** option in the properties dialog of the row.

In view mode, the widgets are compressed and kept in the same row.

5.2.3.9 Group widgets

You can group widgets on the dashboard using the Layout group.

The **Layout group** () consists of two rows and 12 columns. Grouped widgets can be handled as one single widget. You can insert, move, resize, or copy a layout group like any other widget. The **Layout group** is available in the widget bar.

Procedure

1. Insert the  **Layout group** into the dashboard using drag and drop. See Insert widgets in a dashboard (page 30).
2. Resize the layout group. See Resize widgets (page 68).
3. Insert widgets in the layout group.

The layout group is placed on the dashboard.

Use the selection mode if a layout group overlaps its host cells completely and the underlying row and cells cannot be selected. Click **Options** > **Selection mode ON** in the dashboard main menu.

5.2.3.10 Set row style

You can change the style applied to a single row. The selected style specifies, for example, the background color of the row.

Procedure

1. Click a row in the workspace.
You can also select a single row in a layout group
2. Select a style template in the **Style** drop-down menu in the row properties dialog.

The row style is set.

5.2.3.11 Set layout group style

You can change the style applied to a layout group. The selected style specifies, for example, the background color of the layout group.

Procedure

1. Click **Group** inside a layout group.
2. Select a style template in the **Style** drop-down menu in the layout group properties dialog.

The layout group style is set.

5.2.3.12 Set dashboard tab style

You can change the style applied to a dashboard tab. The selected style specifies, for example, the background color of the tab.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click the  **Show menu** icon beside the tab title of the relevant tab.
3. Click the **Style** selection box and select a style.

The dashboard tab style is set.

5.2.3.13 Switch to fixed-grid workspace

You can switch from the smart dashboard workspace to the fixed-grid workspace.

You can only switch if you have not already inserted a widget in the smart dashboard, and you cannot switch from a fixed-grid dashboard to a smart dashboard workspace. If a dashboard was created using the fixed-grid workspace, it is automatically opened in the fixed-grid workspace.

Procedure

1. Create a dashboard (page 11).
2. Click **Manage** > **Switch to fixed grid** in the dashboard main menu.

The fixed-grid workspace is displayed.

5.2.4 Edit a dashboard

You can edit existing dashboards in the dashboard editor.

Depending on the dashboard edit mode used to create the dashboard, the dashboard selected is opened in smart dashboard mode or fixed-grid mode. See Using smart dashboard edit mode (page 11) for details.

Procedure

1. Click **Manage** > **Open** in the dashboard editor main menu.
2. Select an **Available dashboard** and click **OK**. The selected dashboard is opened in edit mode in the dashboard editor.

A warning message is displayed if local changes exist for the dashboard selected. To open the dashboard with the local changes, click **Continue with unsaved local version**.

3. Configure the dashboard.
4. Click **Manage** > **Save** in the dashboard main menu.

Your changes are applied.

To display the dashboard in view mode, click the  **View dashboard** icon in the dashboard main menu.

5.2.5 Delete a dashboard

You can delete dashboards from the dashboard editor.

Warning

Deleted dashboards cannot be restored.

Procedure

1. Open a dashboard to be deleted. (page 15)
2. Click **Manage** > **Delete** in the dashboard main menu.
3. Click **Yes**.

The selected dashboard is deleted.

5.2.6 Edit dashboard properties

You can edit the properties (name, description and tags) of existing dashboards in the dashboard editor.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Manage** > **Properties** in the dashboard main menu. The dashboard **Properties** dialog is displayed.
3. The dashboard name is mandatory and is entered in the **Name** box.
4. The dashboard description is optional and entered in the **Description** box.
5. Comma-separated search tags are optional and are entered in the **Tags** box.
6. Click **OK**.
7. Click **Manage** > **Save** in the dashboard main menu.

Your changes are applied.

5.2.7 Manage dashboard permissions

You can manage dashboard permissions in the dashboard editor. You can assign specific access permissions to individual users or to user groups. If you assign permissions to a user group, the permissions are automatically assigned to all members of that group.

For new users and user groups of a dashboard, you can automatically assign view permissions to all associated assets of the dashboard, such as data feeds and aliases. It is not required to assign the permissions to each asset manually. A user requires the view permission for all associated assets to display the corresponding source data in the dashboard. If view permissions are not assigned to all associated assets, a corresponding option to assign the missing view permissions is additionally displayed in the dialog.

You can assign the following access permissions for saved dashboards only.

- **Edit**

Users can display and edit dashboards in the dashboard editor.

- **View**

Users can view dashboards in the dashboard editor view mode.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Manage > Permissions** in the dashboard main menu. The **Manage dashboard permissions** dialog is displayed.
3. Enter a term in the search box and click **Search**. Clicking **Search** without any input values results in a list of all users and groups.
4. Click **Show MashZone NextGen default groups** to show only default users or user groups in the **Search results** box.
5. Drag a user or a user group from the **Search result** box and drop it on the **Principals with permissions** box.

By default, the creator of the dashboard already exists in the **Principals with permissions** list.

6. Enable or disable the **View** or **Edit** permissions of a user or a user group.
7. Click **Save**.

The button is available if the option **Assign the relevant view permissions to related assets** is disabled, or view permissions are already assigned to all associated assets.

8. Enable the option **Assign the relevant view permissions to related assets** to assign the required view permissions to all associated data feeds and aliases.

The option is available if view permissions are not assigned to all associated assets.

9. Click **Next**.

A new dialog opens. The first list in the dialog contains the assets whose view permissions you can update. The second list contains the assets whose view permissions you cannot change. At least one of the following prerequisites must apply to change the view permissions for data feeds or aliases.

You are an administrator who can edit the permissions for aliases.

You have permissions to view and to edit data feeds.

You have permissions to create and to edit data feeds.

Your changes are applied.

If you want to remove a user or a user group from the **Principals with permissions** list, click the **Delete** icon. Deleted permissions for a dashboard do not affect the associated data feeds or aliases.

5.2.8 Change dashboard style template

You can assign another style template to an available dashboard. Style templates allow you to customize the look and feel of your dashboards, for example, colors schemes, fonts or background color.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Manage** > **Change style template** in the dashboard main menu.
3. Select a style template in the **Dashboard** drop-down menu.
4. Click **Update**.
5. Click **OK**.

The selected style template is applied to the current dashboard.

5.2.9 Use the Process Mining context

The Process Mining context is an easy way to use PPM analyses as data sources for your widgets. With a Process Mining context, you can easily create dashboards based on data from a PPM sever. This allows you to directly access analytics results, such as measures and dimensions, without creating any favorite in PPM itself. In addition, dashboards based on a Process Mining context are automatically filterable, that is, it is no longer required to manually define filters across different widgets.

5.2.9.1 Create a Process Mining context

You can create a Process Mining context (page 18) for a new dashboard or for an already existing dashboard. The Process Mining context is automatically assigned to all new widgets on the dashboard. Widgets that have already been inserted into the dashboard retain their assigned data sources.

Prerequisites

- A PPM connection has been created for each PPM system to be used.
- The appropriate PPM client server must be running to connect to PPM. See the PPM documentation **PPM Installation** for details.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Options** > **Process Mining** > **Create context** in the dashboard main menu. The **Create Process Mining context** wizard opens.
3. Select an available PPM connection from the **PPM connection (alias)** drop-down menu.

The alias is the name of the PPM connection that contains the PPM client connection data defined in ARIS Connect. See Configure PPM server for details. Only PPM connections for which you have the appropriate permissions are displayed.

4. Specify the authentication details to connect to the PPM client server. Select one of the following options.
 - a. Enable the **Single Sign-On** option to log in to the PPM client server via single sign-on (SSO) using your current ARIS Connect credentials.
 - b. Enable HTTP basic authentication and enter the required user name and password of a PPM user.
5. Click **Next**.

Your settings are applied.
6. Select the language in which you want to display the PPM data.

Only the languages specified for the PPM client of the selected PPM connection are available.
7. Select the processes to be analyzed in the **Process type** drop-down menu.

The menu provides the process types specified for the PPM client of the selected PPM connection. Depending on the process type you can select various measures and dimensions.
8. Select the relation to be analyzed from the **Relation** drop-down menu.

The menu provides the relations specified for the PPM client of the selected PPM connection.
9. Select the measures and dimensions to be analyzed.

The measures and dimensions provided depend on the process type you have selected. By default, all available measures and dimensions are preselected.

 - a. Click **Measures** and select the required measures.
 - b. Click **Dimensions** and select the required dimensions.
10. Click **OK** to exit the **Create Process Mining context** wizard.

Your settings are applied. For each selected element, for example, measures and dimensions, the corresponding data columns are created in the Process Mining context. You can now use the context as a data source for your widgets.
11. Click **Next** to manually edit the data columns of the context.
12. Enter a term in the **Search** field to filter the data columns list.
13. If required, you can add a new data column to the Process Mining context.
 - a. Click **Add column**.
 - b. Select the data column type, that is, measure or dimension, from the **Type** drop-down menu.
 - c. Select a measure or dimension from the **Measure** or **Dimension** drop-down menu.
 - d. Enter a column name.
 - e. Click **Add**.

The new data column is added.
14. Click the **Edit** icon to change the settings of a data column. Make your settings and click **Change**.

Besides the type, all other properties can be changed.

15. Click the **Copy** icon to create a new data column based on the copy of the selected one. You must change at least one property. Make your settings and click **Add**.
16. Click the **Delete** icon to delete a data column from the Process Mining context.
17. Click **OK** to exit the **Create Process Mining context** wizard.

Your settings are applied.

The data columns are added to the Process Mining context and can now be assigned as data source columns.

5.2.9.2 Edit a Process Mining context

You can change the settings of the Process Mining context (page 18) configured for a dashboard. Changes in the Process Mining context settings can cause an incorrect data source configuration. Therefore, the data on the dashboard might not be displayed correctly.

If the dashboard uses an older context, you might need to manually migrate the context to the current context version. See Migrate context-based dashboard (page 26).

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Options** > **Process Mining** > **Edit context** in the dashboard main menu.

The last page of the **Create Process Mining context** wizard opens. The page lists all configured data columns of the Process Mining context. See Create a Process Mining context (page 18). By clicking **Previous**, you can navigate to page two and one of the wizard and change the settings.

3. Select an available PPM connection from the **PPM connection (alias)** drop-down menu.
The alias is the name of the PPM connection that contains the PPM client connection data defined in ARIS Connect. See Configure PPM server for details. Only PPM connections for which you have the appropriate permissions are displayed.
4. Enter a term in the **Search** field to filter the data column list.
5. If required, you can add a new data column to the Process Mining context.
 - a. Click **Add column**.
 - b. Select the data column type, that is, measure or dimension, from the **Type** drop-down menu.
 - c. Select a measure or dimension from the **Measure** or **Dimension** drop-down menu.
 - d. Enter a column name.
 - e. Click **Add**.

The new data column is created.

6. Click the **Edit** icon to change the settings of a data column. Make your settings and click **Change**.

You can change all properties except the type.

7. Click the **Copy** icon to create a new data column based on the copy of the selected one. You must change at least one property. Make your settings and click **Add**.
8. Click the **Delete** icon to delete a data column from the Process Mining context.
9. Click **OK** to exit the wizard.

Your settings are applied.

If new data columns are added to the Process Mining context, they can now be assigned as data source columns.

5.2.9.3 Delete a Process Mining context

You can delete a Process Mining context (page 18) configured for a dashboard.

If you delete the Process Mining context of a dashboard, the corresponding data is no longer displayed. The configured context cannot be restored.

Procedure

1. Open a dashboard in the dashboard editor (page 10).
2. Click **Options** > **Process Mining** > **Delete context** in the dashboard main menu.
3. Click **Yes**.

The Process Mining context is deleted.

5.2.9.4 Assign data columns to Process Mining context-based widgets

You can assign the data columns provided by a Process Mining context to widgets.

The Process Mining context is automatically assigned as a data source to all new widgets on the dashboard. You do not need to assign the data source to each widget manually (Assign data (1/2) dialog). Widgets that have already been inserted into the dashboard retain their assigned data sources. The data columns that are provided by the Process Mining context depend on your context configuration (page 18). The procedure to assign data columns is similar to the standard method for assigning data columns to widgets (Seite 31) (Assign data (2/2) dialog). If you do not want to use context-based data columns, you can assign any other data source provided.

Prerequisites

You have created a Process Mining context for the dashboard (page 18).

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click an inserted widget on the dashboard. The relevant properties dialog is displayed.
3. Click  **Assign data** to edit the data source assignment.

The Assign data (2/2) dialog opens.

4. Assign the data columns of the Process Mining context to the widget elements using drag and drop.

For details on the assignment of data columns to specific widgets, see the chapter Assign data sources to widgets (Seite 31) and the following chapters.

5. Click **OK** to save your settings.

The data columns are assigned to the widget elements.

- If required, you can adjust the list of data columns provided for the data columns assignment. Click **Edit context** to edit the data columns specified in the Process Mining context.
- If required, you can change the data source assigned. Click **Use other data** to assign another data source. The **Assign data (1/2)** dialog opens. For details on the assignment of a data source to a widget, see the chapter Assign data sources to widgets (Seite 31). If you assign another data source, you cannot reassign the Process Mining context to the widget.

5.2.9.5 Configure filters for Process Mining context-based widgets

You can configure filter settings for Process Mining context-based widgets.

All widgets on a dashboard that use the Process Mining context are automatically filtered by the selection that you specified for any context-based widget. That is, if you select a data point in a context-based widget, such as a column in a chart, all other context-based widgets on the dashboard are filtered by this selection. The selections are directly passed to PPM and the corresponding filtered data is returned to the dashboard and used as a filter for the relevant widgets.

By default, the filter conditions are also automatically set when you assign data columns of the Process Mining context to the widgets (page 21). Data columns of text and date type assigned to the widget elements are automatically used as filter columns. You do not have to configure the filter conditions manually.

A data column can be used only once in a dashboard as a filter column. Therefore, the filter column is enabled only once on the dashboard to filter across all widgets. By default, this filter column is provided by the widget to which the data column is assigned first. The filter columns used by a widget are listed on the **Context** tab in the properties dialog of the widget. Here you can enable or disable the filter columns that are based on the Process Mining context.

Prerequisites

You have created a Process Mining context for the dashboard (page 18).

Procedure

1. Open a dashboard in the dashboard editor (page 10).
2. Assign data columns to Process Mining context-based widgets (page 21).
3. Click an inserted widget on the dashboard. The relevant properties dialog is displayed.
4. Click **Context** in the properties dialog.

The **Context** tab is available when you have assigned data columns to the elements of the widget. The tab lists all data columns that are assigned to any context-based widget on the dashboard and that can be used as filter columns. The filter columns that can be used by the selected widget are enabled. The filter columns that are used by other widgets are disabled.

5. Enable a filter column for a widget.

You can enable a filter column that is already used by a widget for another widget.

- a. Move the mouse pointer over the  icon to display a tool tip. The tool tip shows the name of the widget that uses the filter column.
- b. Click the name of the widget that uses the filter column. The **Context** tab of the referenced widget opens.
- c. Disable the filter column that you want to reassign.
- d. On the dashboard, click the widget to which you want to assign the filter column.
- e. Click **Context** in the corresponding properties dialog.
- f. Enable the filter column that you want to use for the selected widget.

6. Cancel the selection of a filter column to disable it for filtering, if required.

The filters are configured and can be used in view mode.

You can use the filter criteria configured for context-based widgets in the filter panel (page 24). As long as such a filter criterion is enabled in the filter panel, the filter criterion is no longer available in the widget. In this case you can select the filter values only in the filter panel and not in the widget. For details, see Use the filter panel of Process Mining context-based widgets (page 6).

5.2.9.5.1 Configure a date filter for a Process Mining context

You can configure a filter column of **date** type for a Process Mining context-based dashboard.

You can use the Date filter (page 275) widget to specify a filter column of date type for a Process Mining context-based dashboard. Only one context column of date type can be specified for filtering.

Procedure

1. Click a Date filter widget on the dashboard. The relevant properties dialog is displayed.
2. Click **Context** in the properties dialog to open the tab.
3. Click the drop-down menu and select a filter column.

The menu provides all filterable context columns of date type. **None** is selected if a data column of date type is already used in another widget as a filter column. If you select a filter column from the drop-down menu, the column is used only by the **Date filter** widget.

The filter column is enabled for the **Date filter** widget and can be used in view mode.

5.2.9.6 Exclude a widget from a Process Mining context filter

You can exclude a dashboard widget from being filtered by the Process Mining context.

Prerequisites

You have created a Process Mining context for the dashboard (page 18).

Procedure

1. Open a dashboard in the dashboard editor (page 10).
2. Assign data columns to Process Mining context-based widgets (page 21).
3. Click a widget on the dashboard. The relevant properties dialog is displayed.
4. Click **Context** in the properties dialog.

The **Context** tab is available when you have assigned data columns to the elements of the widget.

5. Enable the **Exclude widget from being filtered** option.

The widget will not be filtered by the Process Mining context. It displays only unfiltered context data and does not respond to context filtering.

The excluded widget can still be used to filter other context widgets.

5.2.9.7 Enable the filter panel

You can enable the filter panel for a Process Mining context-based dashboard.

Prerequisites

You have created a Process Mining context for the dashboard (page 18).

Procedure

1. Open a dashboard in the dashboard editor (page 10).
2. Click **Options > Process Mining > Filter panel** in the dashboard main menu.

The **Filter panel** option is available if you have created a Process Mining context (page 18).

The filter panel is enabled and available in view mode.

You can configure the filter panel in dashboard edit mode. (page 24)

You can configure and use the panel for filtering in dashboard view mode. (page 6)

Warning

If you disable the filter panel, all your settings are lost.

5.2.9.8 Configure the filter panel

You can add and remove filter criteria to the filter panel.

Using the filter panel in view mode (page 6), you can filter widgets on a dashboard whose data is based on a Process Mining context.

- Text-based filter criteria are displayed as a list.
- Numeric filter criteria are provided as a slider in the filter panel. The basic settings are automatically preconfigured, including range, initial selection, step width, and numeric format. You can manually adjust the initial minimum and maximum values of the slider widget. See Slider (**page 290**) widget for details.

Prerequisites

The filter panel is enabled (page 24).

Procedure

1. Open a dashboard in the dashboard editor (page 10).
2. Click the << icon to show the filter panel on the right-hand side of the dashboard.
The panel is hidden by default and no filter columns are specified.
3. To hide the filter panel, click the >> icon.
4. If no filter is specified, add filter criteria to the filter panel.
 - a. Click the  **Add filter** icon to add an initial filter. A corresponding dialog opens.
 - b. Click the  **Filter settings** icon to display filter options.
 - c. Enter a term in the  **Search** box to filter the column list.
The list includes all data columns of text type that have not been assigned as filter columns to the Process Mining context-based widgets. See Assign data columns to Process Mining context-based widgets (page 21).
 - d. Enable the **Selected** option to display only the data columns that have already been selected.
 - e. Select the data columns you want to use as filter criteria.
 - f. Click **OK**.
Your settings are applied. An appropriate filter criterion is added to the filter panel for each selected data column.
5. If the filter panel already contains filter criteria, you can add additional criteria that will be used in the panel.
 - a. Click **Edit filter panel**.
 - b. Click the  **Filter settings** icon to display filter options.
 - c. Enter a term in the  **Search** box to filter the values of a filter criterion.
 - d. Enable the **Selected** option to display only the filter values that have already been selected.
 - e. Select the data columns you want to use as filter criteria.
 - f. You can use the filter criteria configured for context-based widgets (page 22) in the filter panel. As long as such a filter criterion is enabled in the filter panel, the filter criterion is no longer available in the widget. These filter criteria are hidden by default.

1. Enable the option **Show configured context filters from dashboard widgets**. The filter criteria are displayed and marked in bold.
 2. Select the context filter criteria that you want to use in the filter panel.
- g. Click **OK**.
- Your settings are applied. An appropriate filter criterion is added to the filter panel for each selected data column.

The filter panel is configured and the filter values are specified.

You can use the filter panel to filter context-based widgets in dashboard view mode. (page 6)

5.2.9.9 Migrate a dashboard context

With MashZone NextGen version 10.3, the aggregation settings for numeric KPIs in context-based dashboards have been changed.

If you want to edit a dashboard that uses an older context, you must manually migrate the dashboard context to the current context version. The **Context migration** dialog opens, informing you that you must migrate the dashboard context.

During context migration, all numeric columns are automatically deleted and for each PPM KPI the initial column is created again. You must reassign and reconfigure the data columns in the data assignment (**Assign data (2/2)** dialog) for each widget.

Procedure

1. Click **Yes** in the **Context migration** dialog to configure the context settings. The **Edit Process mining context** dialog opens.
2. Edit the context settings (page 20).
3. Click **OK** to save your changes.
4. Software AG recommends that you save your dashboard as a copy to keep the settings of your old dashboard.
 - a. Click **Mange > Save as** in the dashboard main menu.
 - b. Make your settings.
 - c. Click **Save**.
5. Click an inserted widget on the dashboard. The relevant properties dialog is displayed. The **Assign data (2/2)** dialog opens.
6. Reassign the relevant data columns and specify the column options. (page 21)

To specify the KPI value aggregation, click an assigned measure and select an aggregation type in the **Aggregation** drop-down menu.

If you assign a measure column twice and you select different aggregation types, you must change one of the column names.
7. Click **OK** to save your settings.

Your settings are applied and the dashboard context is migrated.

5.2.10 Use Multi-lingual dashboards

With version 10.3 MashZone NextGen supports multi-lingual dashboards. Every widget can offer different elements for translation. This includes for example the widget title, axis title, or partition name. Date format and number patterns are automatically translated by MashZone NextGen. You can also translate certain dashboard elements, such as tab titles.

The dashboard creator can specify different translations in different languages for every widget at creation time (edit mode) and the viewer can then switch between these languages at view time (view mode).

5.2.10.1 Configure languages for translation

You can configure the languages supported by a dashboard. The supported languages can be used to translate individual widget elements.

You can add or remove multiple languages supported by a dashboard. Additionally, you can set the primary design language, which is used by default to create dashboards and widgets.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Manage > Languages** in the dashboard main menu. The corresponding dialog opens.
3. Click **+ Add language**.
4. Select a language in the drop-down menu.
5. To delete a language from the list, move the mouse pointer over a language in the list and click the **Delete** icon. The default language cannot be deleted.
6. To change a language, move the mouse pointer over a language in the list and select a language in the drop-down menu.
7. Click **OK**.

The languages supported by the dashboard are configured.

You can translate the content of the default language in all languages supported by the widget. (page 27)

5.2.10.2 Translate the widget content

You can translate the widget content in all languages supported by a dashboard (page 27). Depending on the widget, different elements can be translated. This includes for example the widget title, axis title, or partition name.

Date format and number patterns are automatically translated by MashZone NextGen.

Prerequisite

You have assigned data columns to the widget. (Seite 31) Except for widgets to which no data can be assigned, such as the **Label** widget.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. The corresponding properties dialog opens.
3. Click  **Translate**.

The **Translations** dialog opens. The default language is the design language and the corresponding content cannot be changed. Only the data source columns that are assigned to widget elements (Seite 31), such as measures, are available for translation.

4. Enter your translation for the content of the default language column in the respective table line.
5. Click **OK**.

The widget content is translated in one or more languages.

You can select the supported languages in view mode.

5.2.10.3 Select the language for your dashboard

You can display your dashboard content, for example, column title, axis title, or axis format, in any language supported by the dashboard.

Prerequisites

The languages supported by the dashboard widgets are configured (page 27) and the widget contents are translated (page 27).

Procedure

1. Display a dashboard in dashboard view mode. (Seite 5)
2. Click **Manage > Languages** in the dashboard main menu and select a language.

The option is available if at least two languages are configured for translation. (page 27)

The dashboard is displayed in the language selected.

5.2.10.4 Translate the tab title

You can translate the tab titles in all languages supported by a dashboard (page 27).

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Add one or more tabs to the dashboard. (page 70)
3. Click a dashboard tab.
4. Click the  **Settings** icon beside the tab title.
5. Click **Translate**.

The **Translations** dialog opens. The default language is the design language and the corresponding content cannot be changed.

6. Enter your translation for the content of the default language in the respective table line.
7. Click **OK**.

The tab title is translated in one or more languages.

You can select the supported languages in view mode.

5.3 Configure widgets

You can configure the visualization and the behavior for all widgets. Additionally, you can assign data sources and set filters for most widgets.

Depending on the widget type, for example, line chart, grid, or image, various options are available.

The appendix lists all available widgets (page 269) and the relevant configurable parameters:

-  Line chart (page 283)
-  Column chart (page 273)
-  Bar chart (page 269)
-  Pie chart (page 287)
-  Bubble chart (page 271)
-  Grid (page 277)
-  Circular gauge chart (page 272)
-  Horizontal and vertical gauge chart (page 279)
-  Traffic light (page 294)
-  Drop-down box (page 276)
-  Input field (page 281)
-  Image (page 280)
-  Label (page 282)
-  Function flow diagram
-  Jump to PPM (page 292)

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. The relevant widget properties dialog is displayed.
3. Click  **Assign data** to edit the data source assignment. See Assign data sources to widgets (Seite 31) for details.
4. Click  **Configure filter** to specify filter conditions. See Define filters for widgets (Seite 60) for details.
5. Click the **Config** tab and set the widget display options. See the list of widgets (page 269) for details.

6. Specify **Actions** for the widget, if required. See Specify actions for widgets for details.
7. Specify **URL selections** for the widget, if required. See Use dynamic URL selection (page 72) for details.

Your settings are applied.

5.3.1 Set up widgets

You can configure the visualization and the behavior in dashboard view mode for all widgets. Additionally, you can assign data sources and set filters for most widgets.

Various options are available depending on the widget type, for example, line chart, grid or image.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. The relevant widget properties dialog is displayed.
3. Click the  **Assign data** icon to edit the data source assignment. See Assign data sources to widgets (Seite 31) for details.
4. Click the  **Configure filter** icon to specify filter conditions. See Define filters for widgets (Seite 60) for details.
5. Click the **Config** tab and set the widget display options. See the list of widget (page 269) for details.
6. Optionally, specify **Actions** for the widget. See Specify actions for widgets for details.
7. Optionally, specify **URL selections** for the widget. See Use dynamic URL selection (page 72) for details.

Your settings are applied.

5.3.2 Insert widgets in a dashboard

Using widgets you can create your individual dashboards and visualize your source data interactively. The dashboard editor provides you with various widgets, for example, line chart, bar chart, speedometer chart, grid, or input field.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget icon in the widget bar. The widget selected is inserted as a blank box on the dashboard and placed in the next empty field. The relevant widget properties dialog is displayed.

Alternatively, place the widget using drag and drop in any empty field on the dashboard. The relevant widget properties dialog is displayed.
3. Click the widget inserted and place it in a empty field on the dashboard using drag and drop.

The widget selected is inserted and placed on the dashboard.

5.3.3 Copy and paste widgets in dashboards

You can copy or cut widgets and paste them in the same or in any other tab of the dashboard.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. A corresponding pop-up menu is displayed.
3. Click the **Copy** icon in the pop-up menu to copy the selected widgets to the clipboard.
4. Click the **Paste** icon in the pop-up menu to insert the copied widgets in the same tab of the dashboard.
5. Click the **Cut** icon in the pop-up menu to cut the selected widget and copy it to the clipboard.
6. Open any tab of the dashboard and press **Ctrl+V**.
Filter relations between copied widgets are retained.

The widget selected is copied to the clipboard and pasted in the selected dashboard tab.

5.3.4 Assign data sources to widgets

Before you can display content in a widget, you must first assign a data source to this widget.

You can define dynamical and reusable input parameters for several data sources, for example, XML, JSON, or ARIS Table. See Create input parameters (page 59). Data sources are not required for **Input field** and **Image** widgets.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. The relevant properties dialog is displayed.
3. Click  **Assign data** to edit the data source assignment.
The **Assign data (1/2)** dialog is displayed. Here you can edit the **Data flow** and create dynamic **Input parameters** for the widget selected. The **Data flow** box shows an overview of the data source configuration.
4. Optionally, you can create input parameters for the selected widget.
Input parameters are dynamic and reusable parameters used in several data source operators and data transformation operators. See Create input parameters (page 59).
5. Select a data source operator in the **Add a data source** bar and specify your settings. See Data source operators (page 296) for a list of available data source operators.
6. In addition to the data source operator you can add further operators to transform the source data. In the **Add data operations** bar, select a data transformation operator, for example, **Change data type**, and specify the settings.

See Data transformation operators (page 317) for a list of available data transformation operators.

7. Click the  **Calculate preview** icon of an operator to display a preview of the transformed source data of the operator. This allows you to track all data changes step by step.

8. Click **Next**.

The **Assign data (2/2)** dialog is displayed.

9. If you want to return to the **Assign data (1/2)** dialog to change the data source settings, click **Previous**.

The data source and transformation operators are assigned to the selected widget.

To display data in a widget (for example, line chart) you must assign the relevant data source columns to the required widget elements. In the **Assign data (2/2)** dialog you can assign data source columns to the individual widget elements. For example, you can assign separate columns to the chart axes as a dimension (X-axis) or KPI (Y-axis).

5.3.4.1 Assign data columns to line, column or bar charts

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

In a line chart (page 283), column chart (page 273), or bar chart (page 269), you can display one dimension (X-axis) and several KPIs (Y-axis), or two dimensions (x-axis and partition) and one KPI (Y-axis).

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click an inserted widget on the dashboard. The relevant properties dialog is displayed.
2. Click  **Assign data** to edit the data source assignment.

The **Assign data (2/2)** dialog is displayed.

3. Assign a **Data column** as a dimension to the **X-axis** using drag and drop.
4. Assign one or more **Data columns** as KPIs to the **Y-axis** using drag and drop.

If you have assigned a column to the **X-axis** and a column to the **Partition**, you cannot assign more than one column to the **Y-axis**.

5. Optionally, assign a **Data column** as a second dimension to the **Partition** using drag and drop.

If you have assigned more than a column to the **Y-axis** the **Partition** is no longer available.

6. Optionally, assign one or more columns to the **More columns (invisible)** element using drag and drop. The columns are not displayed in the widget. They are used for filtering widgets only. See Define filters for widgets (Seite 60) for details.

The data source columns are assigned to the widget elements.

5.3.4.1.1 Set properties of widget elements

For each widget element, such as axis, dimensions or measures, you can edit the settings, for example, axis title, display name or format.

Various settings are available depending on the data type of the assigned data source column. See the Options list below for details.

Procedure

1. Click a source column assigned to the **X-axis** and specify the axis and column settings.
2. Click **Y-axis** to specify the axis settings.
 - a. Click the **Text** tab to specify the **Axis title** and **Axis format**.
 - b. Click the **Data range** tab to limit the KPI value range.
3. Click a source column assigned to the **Y-axis** and specify the settings.
4. Click the source column assigned to the **Partition** and specify the settings.
5. Click a source column assigned to the **More columns (invisible)** box and specify the settings.
6. Click **Coloring** and then click the **Thresholds** tab to specify the KPI thresholds. See Configure KPI thresholds (page 34) for details.
7. Click **Coloring** and then click the **Rating** tab to specify the KPI rating. See Configure KPI rating (page 35) for details.
8. Click **OK** to save your settings and to close the dialog.

The chart is displayed in the dashboard with real data of the assigned data source.

Options list

Option	Description
Axis title	Title of the X- or Y-axis. Enable the title field and enter an Axis title . By default, the Automatic option is selected and the axis title consists of the concatenated names of the assigned columns.
Axis format	Output format of the X- or Y-axis. The selected format overwrites the formats of the assigned columns. By default, Auto is selected and the individual formats are used for tooltips. Optionally, you can add a prefix or postfix to the format, for example, \$ 1,234 or 1,234 mm.
Display name (data point)	New column name displayed in the widget, for example, used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type.
Sorting	Sorts the labels of the X-axis data points in ascending or descending order or it sorts the labels of the X-axis data points

Option	Description
	according to the order of a data column that is assigned to the chart. The option is not available in a line chart if a date column is assigned to the X-axis. The "Sort by" option is not available if a partition is assigned.
Sort by	Sorts the labels of the X-axis data points by any column assigned to the chart. The option is not available if a column is assigned to the Partition .
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Text tab	Enables you to specify the Axis title and Axis format .
Data range tab	Enables you to specify the minimum value (From) and the maximum value (To) of the KPI value range. If the values are not set the values are calculated automatically. You can enter specific values manually or assign data source columns for a dynamic value assignment using drag and drop.
Coloring	Enables you to specify KPI thresholds and ratings. You can set the background or the graphic color depending on the KPI values.
Thresholds tab	Enables you to specify the thresholds values range and to set the background and graphic color accordingly. You must assign at least one KPI to the Y-axis.
Rating tab	Enables you to specify the KPI rating and to set the data point color accordingly. You must assign a single KPI to the Y-axis.

5.3.4.1.2 Configure KPI thresholds

You can define thresholds for numeric and text KPIs to display their status. A colored background and graphic show the threshold range in which a KPI value is located.

Procedure

1. Click **Coloring**.
2. Click **Thresholds** to specify the KPI thresholds.

The tab is only available if at least one column is assigned to the Y-axis.

3. Click **Background** to set the chart background color or click **Graphic** to set the color scheme.

The **Graphic** tab is only available if a single KPI is assigned to the Y-axis.

4. Click a **Color** box to select a background color for each threshold.
5. Select an operator for each threshold to define the KPI value range, for example, < (less than).
6. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric values are valid as dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.

7. Click the **+** **Plus** button to add a threshold or click the **-** **Minus** button to remove a threshold.
8. Select a threshold view style in the **Style** drop-down menu.

This option is only available for the background.

9. Enable the **Colorize axis labels** to apply the specified color scheme to the axis labels.

The thresholds are configured for the selected widget.

5.3.4.1.3 Configure KPI rating

You can define a rating for numeric KPIs to display their status. A colored data point shows the rating range in which a KPI value is located.

For example, you can color the ten poorest or worst KPI values, either as an absolute or as a percentage rating.

Prerequisites

You have assigned a single KPI to the Y-axis.

Procedure

1. Click **Coloring**.
2. Click **Rating** to specify the KPI rating.

The tab is only available if a single column is assigned to the Y-axis.
3. Enable the **Absolute** or **Percentage** option to set an absolute or a percentage rating, for example, to color the top three or 30% KPI values.

The highest value is considered as 100%
4. Click a **Color** box to select a color for each rating range.
5. Select **Top** or **Bottom** in the drop-down menu to assign the selected color to the best respectively the poorest KPI values.
6. Enter a value for the rating range, for example, 5 for the five best KPI values.

7. Click the **+** **Plus** button to add a rating range or click the **-** **Minus** button to remove a rating range.
 8. To color the remaining, unassigned KPI values, enable the **Others** option and select a color.
 9. Enable the **Colorize axis labels** to apply the specified color scheme to the axis labels.
- The rating is configured for the selected widget.

5.3.4.2 Assign data columns to pie charts

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A pie chart (page 287) can display one numeric KPI iterated via a dimension (text or date dimension).

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click an inserted widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a numeric **Data column** to the **KPI** using drag and drop.
4. Assign a text or date **Data column** as a dimension to the **Partition** using drag and drop.
5. Optionally, assign one or more columns to the **More columns (invisible)** element using drag and drop. The columns are not displayed in the widget. They are used for filtering widgets only. See Define filters for widgets (Seite 60) for details.

The data source columns are assigned to the widget elements.

5.3.4.2.1 Set properties of widget elements

You can edit the settings, for example, display name or format, for each widget element, such as measure or dimension.

Various settings are available depending on the data type of the assigned data source column. See the Option list below for details.

1. Click the source column assigned to the **KPI** and specify the settings.
2. Click the source column assigned to the **Partition** and specify the settings.
3. Click a source column assigned to the **More columns (invisible)** box and specify the settings.
4. Click **OK** to save your settings and to close the dialog.

The chart is displayed with real data of the assigned data source.

Options list

Option	Description
Display name (data point)	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, for example, used for data points or tooltip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.

5.3.4.3 Assign data columns to bubble charts

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

In the bubble chart (page 271) you can display one dimension and two KPIs. The two KPIs are plotted on the x- and Y-axis. The dimension is represented by various colors of the individual bubble areas. Optionally, a third KPI can be incorporated; its values determine the size of the bubble areas.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click  **Assign data** to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a **Data column** as a KPI to the **X-axis** using drag and drop.
4. Assign a **Data column** as a KPI to the **Y-axis** using drag and drop.
5. Optionally, assign a **Data column** as a KPI to the **Size by** element using drag and drop.
6. Assign a **Data column** as a dimension to the **Partition** using drag and drop.
7. Optionally, assign one or more columns to the **More columns (invisible)** element using drag and drop. The columns are not displayed in the widget. They are used for filtering widgets only. See Define filters for widgets (Seite 60) for details.

The data source columns are assigned to the widget elements.

5.3.4.3.1 Set properties of widget elements

You can edit the settings, for example, axis title, display name or format, for each widget element, such as axis, dimensions or measures.

Various settings are available depending on the data type of the assigned data source column. See the Options list below for details.

1. Click the **X-axis** or **Y-axis** to specify the axis settings.
 - a. Click the **Text** tab to specify the **Axis title** and **Axis format**.
 - b. Click the **Data range** tab to limit the KPI value range.
2. Click the source column assigned to the **X-axis** and specify the axis and column settings.
3. Click the source column assigned to the **Y-axis** and specify the settings.
4. Click the source column assigned to the **Size by** element and specify the settings.
5. Click the source column assigned to the **Partition** and specify the settings.
6. Click a source column assigned to the **More columns (invisible)** box and specify the settings.
7. Click **Coloring** to specify the KPI thresholds. See Configure KPI thresholds (page 39) for details.
8. Click **OK** to save your settings and to close the dialog.

The chart is displayed with real data of the assigned data source.

Options list

Option	Description
Axis title	Title of the X- or Y-axis. Enable the title field and enter an Axis title . By default, the Automatic option is selected and the axis title consists of the concatenated names of the assigned columns.
Axis format	Output format of X- or Y-axis. The selected format overwrites the formats of the assigned columns. By default, Auto is selected and the individual formats are used for tooltips. Optionally, you can add a prefix or postfix to the format, for example, \$ 1,234 or 1,234 mm.
Display name (data point)	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.

Option	Description
	The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Text tab	Allows you to specify the Axis title and Axis format .
Data range tab	Allows you to specify the minimum value (From) and the maximum value (To) of the KPI value range. If the values are not set the values are calculated automatically. You can enter specific values manually or assign data source columns for a dynamic value assignment using drag and drop.
Coloring > Thresholds	Allows you to specify the thresholds values range and to set the background color accordingly. You must assign at least one KPI to the Y-axis.

5.3.4.3.2 Configure KPI thresholds

You can define thresholds for numeric and text KPIs to display their status. A colored background shows the threshold range in which a KPI value is located.

Procedure

1. Click **Coloring**.
2. Select the axis for which you want to define thresholds.
3. Click a **Color** box to select a background color for each threshold.
4. Select an operator for each threshold to define the KPI value range, for example, < (less than).
5. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric data columns are allowed for dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.

6. Click the **+** **Plus** button to add a threshold or click the **-** **Minus** button to remove a threshold.
7. Select a threshold view style in the **Style** drop-down menu.

This option is only available for the background.

The thresholds are configured for the widget selected.

5.3.4.4 Assign data columns to grids

To display data in the Grid (page 277) widget, you must assign the relevant data source columns to the grid columns required.

You can assign any number of data source columns as grid columns.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign one or more **Data columns** to the **Y-axis** using drag and drop.
4. Optionally, assign one or more columns to the **More columns (invisible)** element using drag and drop. The columns are not displayed in the widget. They are used for filtering widgets only. See Define filters for widgets (Seite 60) for details.

The data source columns are assigned to the widget elements.

5.3.4.4.1 Set properties of widget elements

You can edit the settings for each widget element, for example, column name or format.

Various settings are available depending on the data type of the assigned data source column. See the Options list below for details.

Procedure

1. Click one of the assigned **Grid columns** to specify the settings.
 - a. Click the **Text** tab and specify the settings.
 - b. Click the **Thresholds** tab and specify the KPI thresholds. See Configure thresholds (page 41) for details.
2. Click a source column assigned to the **More columns (invisible)** box and specify the settings.
3. Click **OK** to save your settings and to close the dialog.

The grid is displayed with real data of the assigned data source.

Options list

Option	Description
Axis title	Title of the X- or Y-axis. Enable the title field and enter an Axis title . By default, the Automatic option is selected and the axis title consists of the concatenated names of the assigned columns.

Option	Description
Axis format	Output format of X- or Y-axis. The selected format overwrites the formats of the assigned columns. By default, Auto is selected and the individual formats are used for tool tip. Optionally, you can add a prefix or postfix to the format, for example, \$ 1,234 or 1,234 mm.
New column name	Replaces the data source column name, used for tooltips. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the grid. Available for numeric columns. Enabled by default.
Show cell value	Displays the values of the column. Enabled by default.
Make clickable (for actions)	Makes the column clickable for triggering actions. See Specify actions for widgets for details.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Set as key column for selection	Sets the selected column as key column. When a selection is made and new data arrives, the selection is automatically recreated with the new data. The key column is used to ensure that the new data matches the selection. This means that if values in non-key columns have changed, the selected row remains selected even after new data arrives. You can select one or more columns as key columns.
Text tab	Allows you to specify the New column name , Format and Aggregation .
Thresholds tab	Allows you to specify the thresholds values range.

5.3.4.4.2 Configure thresholds

You can define thresholds for numeric and text grid columns to display the status of a KPI. A colored marker shows the threshold range in which a KPI value is located.

Procedure

1. Click an assigned **Grid columns** to specify the settings.
2. Click the **Thresholds** tab.
3. Click **Background** or **Foreground** to set the KPI values and the background or foreground colors of the grid.

You can define a KPI value range for the thresholds and the corresponding colors of the grid cells background or foreground.

- a. Click a **Color** box to select a color for each threshold.
- b. Select an operator for each threshold to define the KPI value range, for example, < (less than).

The available operators differ for numeric values (for example, < (**less than**) or = (**equals to**)) and text columns (for example, **starts with** or **is equal to**).

- c. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only text data columns are allowed for dynamic threshold values. Assign the required text data column from the data source to a threshold field using drag and drop.

- d. Click the **+** **Plus** button to add a threshold or the **-** **Minus** button to remove a threshold.
- e. Enable the **Colorize row** option to apply the background or foreground color scheme to all cells in the row.

4. Click **Traffic Lights** to set KPI values and to define traffic lights.

You can define a KPI value range for the thresholds and corresponding traffic lights. The traffic lights are displayed in each cell of the rows. By default, there are three traffic light stages and one inactive stage. Additionally, you can select a different shape for each traffic light stage.

- a. Click a **Color** box to select a color and a shape for each traffic light stage.

To use custom images instead of the colored shapes for the traffic light stages, enable the **Image based Traffic Light** option, click the drop-down menu and enter an **Image URL**.

- b. Enable the **Multi State** option to show multiple states of traffic lights instead of a single traffic light.
- c. Select an operator for each threshold to define the KPI value range, for example, is equal to.

- d. Enter a value for each threshold.

You can either enter specific threshold values or you can choose numeric columns from your data source representing the threshold values. Assign a data column from the data source to a threshold field using drag and drop.

- e. Click the **+** **Plus** button to add a traffic light stage or click the **-** **Minus** button to remove a stage.

- f. Enable the **Switch position** option to place the traffic light symbol to the right of the values.
 - g. Enable the **Make clickable (for actions)** option to make the traffic lights clickable for triggering actions. See Specify actions for widgets for details.
 - h. Enable the **Show tooltip** option to show tool tips for each traffic light symbol.
 - i. By default, the traffic light stage is also visible if it is inactive (that is, a KPI value is not available.). Disable the **Inactive state** option if the traffic light stage should not be visible. You can also change the color of the inactive traffic light stage. If you use custom images, you cannot change the color. You can only deactivate the **Show inactive state** option.
5. Click **OK**.

The thresholds are configured for the selected widget.

5.3.4.5 Assign data columns to lists

To display data in the List (page 284) widget, you must assign the relevant data source columns to the widget elements required.

Prerequisite

You have assigned a data source to the widget (Seite 31).

You can assign one or two data source columns as list columns.

Procedure

1. Click a list widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a data column to the List column box using drag and drop.
The list column is the main column by which the list can be sorted and filtered.
4. Optionally, assign a data column to the **Additional column** box using drag and drop.
The additional column shows additional values to the main column.
5. Optionally, assign one or more columns to the **More columns (invisible)** box using drag and drop. The columns are not displayed in the widget. They are used for filtering widgets only. See Define filters for widgets (Seite 60) for details.

The data source columns are assigned to the widget elements.

5.3.4.5.1 Set properties of widget elements

You can edit the settings, for example, column name or format, for each widget element.

Depending on the data type of the assigned data source column, different settings are available. See the **Option list** below for details.

Procedure

1. Click the assigned List column to specify the settings.
2. Click the assigned Additional column to specify the settings.
3. Click an assigned source column of the More columns (invisible) box and specify the settings.
4. Click OK to save your settings and exit the dialog.

The list is displayed in the dashboard with the data of the assigned data source.

Options list

Option	Description
New column name	Replaces the data source column name that is used by default. The column name is also used for the tool tip.
Format	Output format of the column values, for example, used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the grid. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.

5.3.4.6 Assign data columns to circular gauge charts

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A circular gauge chart (page 272) displays a set of aggregated KPI values. The value ranges are arranged in a semicircle with a red pointer that indicates the actual value of the KPI.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a numeric **Data column** to the aggregated **KPI** using drag and drop.

A data source column is assigned to the **KPI** element.

5.3.4.6.1 Set properties of widget elements

You can define a minimum and a maximum value for the displayed KPI value range. Additionally, you can compare the actual KPI value to several threshold values. The value ranges in the speedometer are indicated by various colors.

See the Options list below for details.

1. Click the source column assigned to the **KPI** and specify the settings.
2. Click **Thresholds** to specify the KPI thresholds. See Configure thresholds (page 45) for details.
3. Click **OK** to save your settings and to close the dialog.

The chart is displayed with real data of the assigned data source.

Options list

Option	Description
Display name	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Min value	Minimum value of the displayed KPI value range. If the value is not set the value is calculated automatically.
Max value	Maximum value of the displayed KPI value range. If the value is not set the value is calculated automatically.

5.3.4.6.2 Configure thresholds

You can define thresholds for a numeric KPI to display their status. A colored marker shows the threshold range in which a KPI value is located.

1. Click **Thresholds**.
2. Click a **Color** box to select a color for each threshold.
3. Select an operator for each threshold to define the KPI value range, for example, < (less than).
4. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric data columns are allowed for dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.

5. Click the  **Plus** button to add a threshold or click the  **Minus** button to remove a threshold.

The thresholds are configured for the selected widget.

5.3.4.7 Assign data columns to horizontal and vertical gauge charts

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A horizontal and vertical gauge chart (page 279) displays a set of aggregated KPI values. The value ranges are arranged in a horizontal or vertical bar with a pointer that indicates the actual value of the KPI.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a numeric **Data column** to the aggregated **KPI** using drag and drop.

A data source column is assigned to the **KPI** element.

5.3.4.7.1 Set properties of widget elements

You can define a minimum and a maximum value for the displayed KPI value range. Additionally you can compare the actual KPI value to several threshold values. The value ranges in the speedometer are indicated by various colors.

See the Options list below for details.

1. Click the source column assigned to the **KPI** and specify the settings.
2. Click **Thresholds** to specify the KPI thresholds. See Configure thresholds (page 47) for details.
3. Click **OK** to save your settings and to close the dialog.

The chart is displayed with real data of the assigned data source.

Options list

Option	Description
Display name	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Min value	Minimum value of the displayed KPI value range. If the value is not set the value is calculated automatically.
Max value	Maximum value of the displayed KPI value range. If the value is not set the value is calculated automatically.
Thresholds tab	Allows you to specify the thresholds values range.

5.3.4.7.2 Configure thresholds

You can define thresholds for a numeric KPI to display their status. A colored marker shows the threshold range in which a KPI value is located.

1. Click **Thresholds**.
2. Click a **Color** box to select a color for each threshold.
3. Select an operator for each threshold to define the KPI value range, for example, < (less than).
4. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric data columns are allowed for dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.

5. Click the  **Plus** button to add a threshold or the click  **Minus** button to remove a threshold.

The thresholds are configured for the selected widget.

5.3.4.8 Assign data columns to traffic lights

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A traffic light (page 294) displays the status of a KPI. A colored marker shows the threshold range in which a KPI value is located.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a numeric or text **Data column** to the **KPI** element.

A data source column is assigned to the **KPI** element.

5.3.4.8.1 Set properties of widget elements

For the displayed KPI value range you can define a minimum and a maximum value. Additionally you can compare the actual KPI value against several threshold values. The value ranges in the speedometer are indicated by various colors.

See the Options list below for details.

Procedure

1. Click the source column assigned to the **KPI** and specify the settings.
2. Click **Thresholds** to specify the KPI thresholds. See Configure thresholds (page 49) for details.
3. Click **OK** to save your settings and to close the dialog.

The traffic light is displayed using the real data of the assigned data source.

Options list

Option	Description
Display name	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.

Option	Description
	The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.

5.3.4.8.2 Configure thresholds

You can define thresholds for numeric and text columns to display the status of a KPI. A colored marker shows the threshold range in which a KPI value is located.

By default, there are three traffic light stages plus the inactive stage.

Procedure

1. Click **Thresholds**.

The option is only available if a column has been assigned to the KPI.

2. Select **Color and shape** in the drop-down menu to display the traffic light stages with various colored shapes.
3. Click a color box to select a color and a shape for each threshold stage.
4. Select **Image** in the drop-down menu to use custom images instead of the colored shapes for each traffic light stage.
5. Select an operator for each threshold to define the KPI value range, for example, < (less than).

The available operators differ for numeric values (for example, < (**less than**) or = (**equals to**)) and text columns (for example, **starts with** or **is equal to**).

6. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric data columns are allowed for dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.

7. Click the **+** **Plus** button to add a threshold or the **-** **Minus** button to remove a threshold.
8. By default, the traffic light stage is also visible if it is inactive (A KPI value is not available.). Disable the **Inactive state** option if the traffic light stage should not be visible. You can also change the color of the inactive traffic light stage. If you use custom images, you cannot change the color. You can only deactivate the **Show inactive state** option.

The thresholds are configured for the selected widget.

5.3.4.9 Assign data columns to a drop-down box

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A drop-down box (page 276) provides you with a selection of values in a drop-down menu for you to filter other widgets.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a **Data column** to the **Visible column** element using drag and drop.
4. Optionally, assign one or more columns to the **More columns (invisible)** element using drag and drop. The columns are not displayed in the widget. They are used for filtering widgets only. See Define filters for widgets (Seite 60) for details.

The data source columns are assigned to the widget elements.

5.3.4.9.1 Set properties of widget elements

For each widget element you can edit the settings, for example, display name or format.

See the Options list below for details.

1. Click the source column assigned to the **Visible column** element and specify the settings.
2. Click a source column assigned to the **More columns (invisible)** box and specify the settings.
3. Click **OK** to save your settings and to close the dialog.

The drop-down menu is displayed with real data of the assigned data source.

Options list

Option	Description
Display name	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.

5.3.4.10 Assign data columns to labels

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A label (page 282) displays a fixed text you have entered, or a text that is supplied dynamically by a data source.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a source **Data column** to the **Data column** element using drag and drop.

A data source column is assigned to the widget element.

If you have assigned a text or date column to the widget the values are retrieved from the first data row.

5.3.4.10.1 Set properties of widget elements

You can edit the settings, for example, display name or format for each widget element.

See the Options list below for details.

Procedure

1. Click the source column assigned to the **Visible column** element and specify the settings.
2. Click a source column assigned to the **More columns (invisible)** box and specify the settings.
3. Click **OK**.

Your settings are applied. The drop-down menu in the dashboard is displayed with real data of the assigned data source.

Options list

Option	Description
Display name	New column name displayed in the widget, that is for instance used for KPI, data point or tool tip. By default, the data source column name is used.
Format	Output format of the column values, that is for instance used for data points or tool tip. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Initial selected	Specifies the initial selected column value. By default, the first column value

Option	Description
value	is selected.
Sort descending	Displays the column values in a descending order in the drop-down menu.

5.3.4.11 Assign data columns to a rich text area

To display data in the Rich text area (page 288) widget, you must assign the relevant data source columns to the required widget elements.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a Rich text area (page 288) widget in the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign one or more data source columns to the **Data column** box using drag and drop.

The data source columns are assigned to the widget elements.

5.3.4.11.1 Set properties of widget elements

You can edit the settings, for example, the display name or the format for each widget element. Settings differ depending on the data type of the assigned data source column. See the option list below for details.

Procedure

1. Click an assigned column in the **Data column** box to specify the settings.
2. Click **OK**.

Your settings are applied. The drop-down menu in the dashboard is displayed with data of the assigned data source.

Options list

Option	Description
New column name	Replaces the default name of the data source column. The column name is also used for the tool tip.
Format	Output format of the column values, that is used for data points or tool tips. Available for columns of date and numeric type.
Round	Displays rounded KPI values. Available for numeric columns. Enabled by

Option	Description
numerically	default.

5.3.4.12 Assign data columns to sliders

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

The data assigned are used to define a data range for the Slider widget. You can assign numeric data source columns used as minimum and maximum of the data range. Optionally, you can assign columns used as default values. If you do not assign a column to a widget element, you can manually specify the required value in the widget properties. For details, see the Slider widget.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a Slider widget in the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a data column to the **From** and **To** boxes under **Range** using drag and drop.
The values are used as minimum or maximum values of the data range.
4. Assign a data column to the **From** and **To** boxes under **Default selection** using drag and drop.
The values are used as initial minimum or maximum values of the data range.

The data source columns are assigned to the widget elements.

5.3.4.12.1 Set properties of widget elements

You can edit the settings, for example, column name or aggregation, for each widget element.

Procedure

1. Click an assigned data column.
2. Specify your settings.
3. Click **OK** to save your settings and exit the dialog.

The widget is displayed with data of the assigned data source columns.

Option list

Option	Description
--------	-------------

Display name	Replaces the data source column name that is used by default. The display name is also used for the tool tip.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.

5.3.4.13 Assign data columns to date filters

To display data in the widget, you must assign the relevant data source columns to the widget elements.

The data assigned are used to define a date range for the Date filter (page 275) widget. You can assign data source columns of **Date** type that are used as start and end values of the date range. Optionally, you can assign columns used as default values. If you do not assign a column to a widget element, you can manually specify the required value in the widget properties. See Date filter (page 275) for details.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a Date filter (page 275) widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a data column to the **From** and **To** boxes under **Range** using drag and drop.
The values are initially used as start or end values of the date range.
4. Assign a data column to the **From** and **To** boxes under **Default selection** using drag and drop.
The values are used as default start or end values of the initially selected data range.
5. Click **OK** to save your settings and exit the dialog.

The data source columns are assigned to the widget elements.

5.3.4.14 Assign data columns to vector map

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A vector map (page 295) allows you to interact with a vector-based online map and to visualize geographic areas of interest.

Even without any data assignments, based on the template chosen, you can select regions on the map. This allows you to trigger selection events and to filter other widgets.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a **Vector map** widget on the dashboard. The relevant properties dialog is displayed.
2. Click  **Assign data** to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a data column containing region identifiers to the **ID** element using drag and drop.
4. Assign a data column to the **Region value** element using drag and drop. The element provides numeric values for the individual regions.
5. Assign a data column to the **Region name** element using drag and drop. The element provides names for the individual regions, that are used for the labels & tooltips.

The data source columns are assigned to the widget elements.

5.3.4.14.1 Set properties of widget elements

For each widget element, such as axis, dimensions or measures, you can edit the settings, for example, axis title, display name or format.

Various settings are available depending on the data type of the assigned data source column. See the Options list below for details.

Procedure

1. Click the source column assigned to the **ID** and specify the display name and the initial selected value.
2. Click the source column assigned to the **Region value**.
 - a. Click the **Text** tab and make your settings..
 - b. Click the **Thresholds tab** to specify the KPI thresholds. See Configure KPI thresholds (page 56) for details
3. Click the source column assigned to the **Region name** and make your settings.
4. Click **OK** to save your settings and to close the dialog.

The chart is displayed in the dashboard with real data of the assigned data source.

Options list

Option	Description
Display name (data point)	New column name displayed in the widget, for example, used for KPI, data points or tool tips. By default, the data source column name is used.
Initial selected value	Initial selected ID. If an ID is specified, the corresponding region

Option	Description
	will be selected by default.
Text tab	Enables you to specify the column settings.
Thresholds tab	Enables you to specify the thresholds values range and to set the background and graphic color accordingly.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Format	Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.

5.3.4.14.2 Configure KPI thresholds

You can define thresholds for numeric KPIs to display their status. A colored background shows the threshold range in which a KPI value is located.

Procedure

1. Click a **Color** box to select a background color for each threshold.
2. Select an operator for each threshold to define the KPI value range, for example, < (less than).
3. Enter a value for each threshold.
You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric values are valid as dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.
4. Click the **+** **Plus** button to add a threshold or click the **-** **Minus** button to remove a threshold.

The thresholds are configured for the selected widget.

5.3.4.15 Assign data columns to map with markers

To display data in a widget, you must assign the relevant data source columns to the required widget elements (for example, to the X-axis).

A map with markers (page 286) allows you to mark points of interest on a map based on geo-coordinates (latitude & longitude) defined in the assigned data.

Prerequisite

You have assigned a data source to the widget (Seite 31).

Procedure

1. Click a **Map with markers** widget on the dashboard. The relevant properties dialog is displayed.
2. Click  **Assign data** to edit the data source assignment.
The **Assign data (2/2)** dialog is displayed.
3. Assign a data column containing marker identifiers to the **ID** element using drag and drop.
4. Assign a data column of numeric type to the **Latitude** element using drag and drop. The assigned column should contain the required latitudes for the markers.
5. Assign a data column of numeric type to the **Longitude** element using drag and drop. The assigned column should contain the required longitudes for the markers.
6. Assign a data column of numeric, or text type to the **Marker value** element using drag and drop. The element provides numeric or text values for the corresponding markers.
7. Assign a data column of numeric, text, or date type to the **Marker label** element using drag and drop. The element provides labels for the corresponding markers.

The data source columns are assigned to the widget elements.

5.3.4.15.1 Set properties of widget elements

For each widget element, such as axis, dimensions or measures, you can edit the settings, for example, axis title, display name or format.

Various settings are available depending on the data type of the assigned data source column. See the Options list below for details.

Procedure

1. Click the source column assigned to the **ID** and specify the display name and the initial selected value.
2. Click the source column assigned to the **Latitude** and specify the display name.
3. Click the source column assigned to the **Longitude** and specify the display name.
4. Click the source column assigned to the **Marker value**.
 - a. Click the **Text** tab and make your settings.
 - b. Click the **Thresholds tab** to specify the KPI thresholds. See Configure KPI thresholds (page 58) for details
5. Click the source column assigned to the **Marker label** make your settings.
6. Click **OK** to save your settings and to close the dialog.

The chart is displayed in the dashboard with real data of the assigned data source.

Options list

Option	Description
Display name (data point)	New column name displayed in the widget, for example, used for KPI, data points or tool tips. By default, the data source column name is used.
Initial selected value	Initial selected ID. If an ID is specified, the corresponding region will be selected by default.
Text tab	Enables you to specify the column settings.
Thresholds tab	Enables you to specify the thresholds values range and to set the background and graphic color. You must assign at least one KPI to the Y-axis.
Aggregation	Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value. The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.
Format	Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.

5.3.4.15.2 Configure KPI thresholds

You can define thresholds for numeric KPIs to display their status. A colored background shows the threshold range in which a KPI value is located.

Procedure

1. Click a **Color** box to select a background color for each threshold.
2. Select an operator for each threshold to define the KPI value range, for example, < (less than).
3. Enter a value for each threshold.

You can either enter specific threshold values or you can choose columns from your data source representing dynamic threshold values. Only numeric values are valid as dynamic threshold values. Assign the required numeric data column from the data source to a threshold field using drag and drop.

4. Click the  **Plus** button to add a threshold or click the  **Minus** button to remove a threshold.

The thresholds are configured for the selected widget.

5.3.4.16 Create input parameters

You can create input parameters to enter dynamic values that are passed to the data transformation step (for example, for filtering) or passed to the data source itself.

You can use parameter sets to create dynamic URLs for your XML source and RAQL queries. The parameters are used to dynamically pass context from dashboards to applications using the URL. You can create dynamic URLs and query parameters that provide the required context to open URLs, to invoke web services, etc.

Procedure

1. Click a widget on the dashboard. The relevant properties dialog is displayed.
2. Click the  **Assign data** icon to edit the data source assignment.
The **Assign data (1/2)** dialog is displayed.
3. Click the **Input parameters** box.
4. Click the input parameter type required in the **Add input parameters** bar.
The data types **Text**, **Number** and **Date** are available for input parameters. They allow the dynamic entry of text, date or numerical values in data feed processing.
5. Enter a **Name** for your parameter set.
6. Enter an optional **Default value**.
The default value is used if no value is passed from the dashboard to the input parameter.
7. Enter a **Preview value**.
The preview value is used to calculate a data preview (debug run) in the data assignment dialog. See Assign data sources to widgets (Seite 31).
8. Click the **Data flow** box to display the data source assignment overview.

The input parameter is displayed in the **Input parameters** box. The defined input parameters are only available for the currently selected widget.

You can use the input parameters for the **XML**, **CSV**, **MS Excel**, and **ARIS Table** data sources. Available input parameters are provided by the data source operator via an additional button.

If you use the input parameter in an XML data source URL, you can specify that the value of the parameter must be URL encoded. For an input parameter of **Date** or **Number** type, you can additionally specify the format pattern in the data source operator.

To specify the parameter properties click the inserted input parameter and select the option required.

Input parameters can also be used as filter elements for several data transformation operators, for example, Insert column, to enter dynamic values in columns. The list of parameters provided, depends on the selected data type of the relevant column.

The input parameters are also provided as filter elements by the widget. See Define filters for widgets (Seite 60) for details.

5.3.5 Display legend

You can display a legend in a widget.

A legend indicates the assignment of the colors in the chart to individual elements in the widget (for example, dimensions). Most widgets provide the option to display a legend on the side of the widget. If there are too many legend entries then legends can be navigated using pagination arrows. The arrows can only be used in view mode.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click on a widget on the dashboard. A corresponding pop-up menu is displayed.
3. Show the **Config** tab to set the display options.
4. Select the position where the legend is displayed in the widget. Default is **None**, i.e., no legend is display.

The legend is displayed on the selected position in the widget.

5.3.6 Remove widgets from dashboard

You can remove widgets from a dashboard.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Select one or more widgets on the dashboard. A corresponding pop-up menu is displayed.
3. Click the **Delete** icon in the pop-up menu.
4. Click **Delete**.

The widgets selected are deleted from the dashboard.

5.3.7 Define filters for widgets

In dashboard edit mode, you can define relations between widgets by specifying filter conditions for selected widgets. The defined filters can be used in dashboard view mode (See Use interactive filters in your dashboards (Seite 5)).

Most widgets support data filters. Especially, the **Drop-down box** (page 276), **Input field** (page 281) and **List** (page 284) widgets can be used for filtering other widgets. (See Insert widgets to a dashboard (page 30)).

FILTERING USING MULTIPLE SELECTION

Multiple selection allows you to filter widgets by selecting multiple values, such as multiple rows or data points, in a widget. See Use interactive filters in dashboards (Seite 5) for details.

In particular, the List (page 284) widget is provided for filtering using multiple selection.

The values selected in a widget are processed as a list and passed on to the widget to be filtered. The widget to be filtered must also support multiple selection. This means that the filtered widget is enabled to process a list of values

and not only a single value. You can enable a widget and, respectively, the assigned data feed to receive and process a value list by inserting user inputs (input parameters (page 59)) of **List** type in the data feed, Date user input (List) (page 358), Number user input (List) (page 359), and Text user input (List) (page 361).

If user inputs (input parameters) of List type have been defined, they are listed below the **Source columns** of a widget. They are identified by their own list icon.

Note that user inputs (input parameters) of **List** type cannot be used to filter single values. For details see Configure filters for Process Mining context-based widgets (page 22) and Use the filter panel of Process Mining context-based widgets (page 6).

PROCESS MINING CONTEXT-BASED DASHBOARDS

The filter definition method described below is not supported by Process Mining context based widgets. The filter conditions are automatically set when you assign data columns of the Process Mining context to the widgets.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. If you have placed a widget on several tabs, select the tab containing the widget for which you want to set a filter. You cannot change the tab in the filter dialog. See Display widgets on several dashboard views (page 69) for details.
3. Click a widget on the dashboard. The relevant properties dialog is displayed.
4. Click the  **Filter** icon to edit the filter configuration. The filter dialog of the selected widget displays the **Source columns** assigned to the widget and the **Select filter widget** bar. See Assign data sources to widgets (Seite 31) for details.

If applicable, input parameters are displayed as additional filter widgets. If input parameters have been defined, they are listed below the **Source columns** of a widget. See Create input parameters (page 59) for details.

If invisible columns have been defined, they are listed as **Source columns**. Invisible columns are not displayed in the widget, but they can be used as a filter criterion. See Assign data sources to widgets (Seite 31) for details.

5. Drag the column that you want to filter from the **Source columns** box and drop it into the **Drop here to define a new filter** field. The **Define filter condition** dialog is displayed. All widgets that can be used as a filter are highlighted. The columns in the **Available coordinates** box for each widget can be used as filter condition values.
6. Define the filter condition.

You can use a column or a constant value as filter condition.

- a. Select a condition from the drop-down menu, for example, **is equal to**.
- b. Drag a coordinate from the **Available coordinates** box of the widget that you want to use as filter widget and drop the selected coordinate into the empty filter condition field.

- c. Alternatively, you can enter a constant value as filter condition.
 - d. To enable case sensitivity, click the **Aa Match case / Ignore case** icon. This option is available only for coordinates or values of type **Text**.
 - e. If the filter conditions are considered to be fulfilled, even if the selected values are empty, click the  **Empty compare values are accepted** icon.
 - f. Click the **Add** icon to add a filter condition.
 - g. Enable the **block values** option to block the rows that meets the condition.
7. Click **Add an additional filter** to add filters to other source columns.
 8. Click **Save filter**.

Your filter conditions are saved and can be applied to your dashboards (Seite 5).

5.3.8 Specify actions for widgets

You can assign actions to specific widgets (for example, charts, traffic light, label, image). The actions either select and apply data in widgets, or open an other dashboard tab, or open a specific URL.

You can specify a data selection for a specific widget on a dashboard tab, for example, a column in a table. The widget is then displayed with the selected data applied. If the selected data also serves as a filter for other widgets, these widgets are filtered accordingly.

Depending on the selected widget, the actions are triggered by clicking a widget, by a mouse over event or by a selection change event.

On mouse over events are performed if you move your mouse pointer over a data point of a widget, for example, a coordinate of a line chart. **On selection change** events are performed if you click a data point in a chart or if you delete a data selection in a widget.

For the **Grid** widget additional trigger options can be available, depending on your **Grid** widget settings. That are **On "item" click** and **On "item" traffic light click**. Whereby "item" is a placeholder for a column name. In the Grid widget settings you can make column cells clickable for triggering actions. See Assign data columns to grids (page 40) for details.

You can use the defined action in the dashboard view mode. See Open a dashboard in MashZone NextGen (Seite 4). Depending on the widget, you can delete a selection by clicking in the background or selecting the **Clear selection** option in the pop-up menu () of the widget.

5.3.8.1 Change tab

By triggering, the action calls another dashboard tab.

The action is only available if you have added one or more views in the dashboard. See Add dashboard tabs (page 70).

If you deactivate the option, the action will not be deleted but deactivated in view mode.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a chart widget on the dashboard that supports actions. The relevant properties dialog is displayed.
3. Click the **Action** tab.
4. Select the action trigger event in the **Trigger** drop-down menu.
The actions available dependent on the widget selected.
5. Activate the **Change** tab option.
6. Select the target tab in the drop-down menu.

Your settings are applied and your action is specified for the selected widget.

5.3.8.2 Set selection

By triggering, the action sets a data selection in a target widget. The target widget is displayed applying the data selected.

This action sets a specific selection, for example, a column in a table or a data point in a chart, in one or several target widgets. The target widgets can be placed on any tab available in the dashboard. If the data selected also represents filter values for another widget, this widget is filtered accordingly.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a chart widget on the dashboard that supports actions. The relevant properties dialog is displayed.
3. Click the **Action** tab.
4. Select the action trigger event in the **Trigger** drop-down menu.
The actions available dependent on the widget selected.
5. Activate the **Set selection** option.
If you deactivate the option the action will not be deleted but deactivated in view mode.
6. Configure the **Set selection** action.
 - a. Click **Configure** to set the action configuration.
Initially all **Available coordinates** of the target widgets are displayed.
 - b. Set the coordinates of the widget you want to select by your action. For this, drag the relevant coordinate and drop it in the field of the **Select selection** widget area.
 - c. Set the values that should be selected. You can enter a constant value or you can assign the values of a coordinate of another widget. Select a tab, drag the relevant widget coordinate, and drop it in the **Selection** box of the previously set coordinate.
Only the coordinates with the fitting values are provided.
 - d. Click **Add an additional selection** to define a further selection.

- e. Click **Save action**.

Your settings are applied and your action is specified for the selected widget.

5.3.8.3 Call URL

By triggering, the action calls a specific URL.

The target widgets can be placed on any tab available in the dashboard.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a chart widget on the dashboard. The relevant properties dialog is displayed.
3. Click the **Action** tab.
4. Select the action trigger event in the **Trigger** drop-down menu. The actions available dependent on the widget selected.
5. Enable the **Call URL** option.

If you disable the option, the action is not deleted but it is deactivated in view mode.

6. Set the action configuration.

- a. Click **Configure**.

- b. Enter the target URL in the **URL** input box of the **Enter target URL** area.

You have the option to create a dynamic URL by adding available coordinates to the URL. For this you can insert the coordinates from several widgets via drag and drop. A selected coordinate will be placed on the current cursor position in the **URL** input field.

You can also add coordinates of widgets placed on any tab available in the dashboard. In this case select a tab first and then insert the relevant **Available coordinates** in the **URL** input field via drag and drop.

In case of a coordinate of type **number** or **date**, click the inserted coordinate and select a format pattern.

To ensure that a coordinate is URL encoded, click the inserted coordinate and activate the **Use URL encoding** option.

- c. In the Target window field you can enter a name of the window where the URL should be opened, or you can select a target attribute in the drop-down menu. Available attributes are **_blank** (new window), **_self** (self window), **_parent** (parent window), and **_top** (entire window).
- d. Click **Save action**.

Your settings are applied and your action is specified for the selected widget.

5.3.8.4 Specify an image source URL

You can configure the source of an image (page 280) to be displayed on a dashboard. You can specify a specific URL or select a URL alias as image source. You can also define a dynamic URL to use in view mode.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click an inserted image on the dashboard. The relevant properties dialog is displayed.
3. Click the  **Source URL** icon. The corresponding dialog opens.
4. To define a specific URL, select **URL** from the **Source** drop-down menu and enter a URL in the URL input field.
5. To select a URL alias, select **URL alias** from the **Source** drop-down menu and select a alias from the **Path prefix (alias)** drop-down menu.

The **Path prefix (alias)** drop-down menu provides aliases configured in the administration. Only URL aliases without basic authentication are provided. Only the URL aliases for which you have the appropriate permissions (page 200) are provided.

6. Enter a URL in the **Default source** input box to specify an image to be displayed when no selection is made. The image is also displayed in edit mode. If this box is left empty no image is shown.
7. Define a dynamic URL using selections (coordinates). The dynamic URL takes effect only in view mode.
 - a. You can create a dynamic URL by adding available coordinates to the URL. For this you can insert the coordinates from several widgets using drag and drop. A selected coordinate is placed on the current cursor position in the URL input field.
 - b. You can also add coordinates of widgets placed on any tab available in the dashboard. In this case, select a tab first and then insert the relevant available coordinates in the URL input field using drag and drop.
 - c. Click **Configure input** to change the format of a coordinate to type **Date** or **Number**. The option is enabled only if at least one coordinate is inserted in the URL input field.
 - d. In case of a coordinate of type **number** or **date**, click the inserted coordinate and select a format pattern.
 - e. To ensure that a coordinate is URL encoded, click the inserted coordinate and activate the **Use URL encoding** option.

8. Click **OK**.

Your settings are applied and the URL is specified for the selected image.

See also

Image (page 280) widget

5.3.8.5 Post data

The action creates an outbound API to pass data from MashZone NextGen dashboards to an embedding system, for example, an external web application.

The API data structure consists of the dashboard id, the specified external widget id (**URL-ID**, see: Use dynamic URL selection (page 72)), the selected coordinate names (columns), the selected coordinate values, and the name of the event triggering the outbound data push. See the **MZNG outbound data structure** example below.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click an inserted chart widget that supports actions. The relevant properties dialog is displayed.
3. Click the **Action** tab.
4. Select the action trigger event in the **Trigger** drop-down menu.

The available actions depend on the selected widget.

5. Activate the **Post data** option.

If you disable the option, the action is not deleted but it is deactivated in view mode.

6. Click **Configure** to set the action configuration.
 - a. Select the coordinates that you want to add to the outbound data.

The coordinates selected are added to the outgoing data, that is displayed in the Outbound data preview box.
 - b. Click **Save action**.

Your settings are applied and your action is specified for the selected widget.

Example

The outbound data is structured as follows:

```
{
  "dashboardGUID": "d216bf4a-bd12-476d-aa5d-2a07a3efd4bf",
  "outbound Widgets": [
    {
      "extId": "widget2",
      "outboundData": [
        { "name": "ARTIST",
          "value": "Bob Dylan",
          "type": "TEXT"
        },
        {
          "name": "PRICE",
          "value": "11.0",
          "type": "NUMERIC"
        }
      ]
    },
    {
      "trigger": "onSelectionChange"
    }
  ]
}
```

By triggering the **Post data** action, MashZone NextGen sends the configured outbound data using **window.postMessage()** events. In order to receive the events in an embedding system a listener function must be implemented as in the example below:

```
function listener(event){
// The origin of the window that sent the message
// at the time postMessage was called
// Format: protocol://host:port var origin = event.origin
// A reference to the window object that sent the message var source = event.source;
// The posted data object var data = event.data;
}
if (window.addEventListener){ addEventListener("message", listener, false) } else
{ attachEvent("onmessage", listener) }
```

5.3.8.6 Send events

Triggering the **Sends event** action, MashZone NextGen sends an event to a defined target. The action is triggered by a data point selection.

Before you can use this action you must define a target for receiving events. From MashZone NextGen version 10.1 Apama is supported as event receiving target. You can use an Apama event target specified in MashZone NextGen admin console for sending events to a running Apama system. See Manage Apama Event Targets for details.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click an inserted chart widget that supports actions. The relevant properties dialog is displayed.
3. Click the **Action** tab.
4. Select the action trigger event in the **Trigger** drop-down menu.

The available actions depend on the selected widget.

5. Activate the **Send event** option.

If you disable the option, the action is not deleted but it is deactivated in view mode.

6. Set the action configuration.

- a. Click **Configure**.
- b. Select an Apama event target in the **Apama event target** drop-down menu.

An Apama event target specifies a combination of Apama system and Apama event type that can be used to send events from MashZone NextGen to Apama.

By selecting a Apama event target several boxes are displayed representing fields of the event type configured in the Apama event target. The event fields can be used as container for data that is sent by MashZone NextGen.

- c. Enter a fixed value manually into an event field or insert a coordinate that you want to be sent. You can insert a coordinate via drag and drop.

You can also insert coordinates of widgets placed on any tab available in the dashboard. In this case select a tab first and then insert the relevant coordinates in the target event fields via drag and drop.

- d. Click **Save action**.

Your settings are applied.

5.3.9 Hide or display widget header and border

You can hide the header and the outline of the widget container.

The header and the outline of certain widgets are hidden by default, for example, Input field and Combobox.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. A corresponding pop-up menu is displayed.
3. Click the **Config** tab to set the display options.
4. Click **More options**.
5. Set the **Container** header and border.
 - Click the **Hide header** icon. This option hides the header and the title of the widget and resizes the content. To unhide the header, click the icon again.
 - Click the **Hide border** icon. This option hides the outline of the widget container. To unhide the outline, click the icon again.

The header and outline of the widget are hidden or displayed.

5.3.10 Resize widgets

You can scale the size of widgets.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. The widget is displayed with a frame.
3. In smart dashboards
 - a. Resize the widget width by dragging the anchor point of the frame with your mouse pointer.

A widget width can be resized across multiple empty fields.
 - b. Resize the widget height by resizing the height of the row in which the widget is inserted. To resize the row height, drag the upper or lower row border with your mouse pointer.

The height of all widgets that are inserted into the same row is resized automatically.
4. In fixed-grid dashboards drag the anchor point of the frame with your mouse pointer.

The selected widgets are resized.

5.3.11 Move widgets to the foreground or background

On the dashboard, you can move widgets to the foreground or background. For example, you can display a chart in the background of the dashboard and place several widgets on top of it.

This option is only available in fixed-grid dashboard mode. See [Switch to fixed-grid workspace](#) (page 15) for details.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Select one or more inserted widgets in the dashboard. A corresponding pop-up menu is displayed.
3. Set **Layering**
 - Click **Bring to front**. Displays the widget in front of one or more other widgets.
 - Click **Bring forward**. Brings the widget one level forward.
 - Click **Send backward**. Brings the widget one level backward.
 - Click **Send back**. Displays the widget behind one or more other widgets.

The selected widgets are moved forward or backward on the dashboard.

5.3.12 Display widgets on multiple dashboard tabs

You can display a widget on multiple views of a dashboard.

A widget that is placed on multiple tabs can be used for filtering other widgets and for triggering selection events across multiple tabs. At the same time, the global placed widget can be filtered and triggered by other widgets. If a global widget is selected (for example, if you click a data point of a line chart), the selection is shown on all tabs on which the widget is placed. The dependent filters and actions are triggered for the entire dashboard.

The option is only available for dashboards that use the fixed grid work space.

On each tab, the widget has the same features such as position, size, configuration and filter conditions etc., except for the widget layering. See [Move widgets to the foreground or background](#) (page 69).

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget. The relevant properties dialog is displayed.
3. Click **More options**.
4. Click **Select tabs** and select the tabs on which the widget should be placed.

Before you can display a widget on several views, you must first add one or more tabs. See [Add dashboard tabs](#) (page 70) for details.
5. To remove a widget from a tab, deselect the relevant tab in the **Select tabs** menu.

You cannot remove a widget from the currently activated tab.
6. To delete a widget click **Delete** in the pop-up menu.

If you delete a widget on a tab, the widget is deleted on all tabs on which it is used. The widget is placed on several tabs and displayed on the corresponding dashboard tabs. You can use an action defined in dashboard view mode. See Use dashboards in view mode (page 5).

5.3.13 Change widget style template

You can assign an other style template to a widget on your dashboard. Style templates let you customize the look and feel of your widgets, for example, colors schemes, fonts or background color.

The style templates provided for the selected widget are part of the style template that is assigned to the current dashboard (page 18).

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget on the dashboard. The relevant properties dialog is displayed.
3. Click **More options**.
4. Select a style template in the **Style** drop-down menu.

The selected style template is applied to the current widget.

5.4 Configure dashboard tabs

You can configure the individual dashboard tabs in the dashboard editor.

5.4.1 Add dashboard tabs

You can add additional dashboard tabs. The individual views are displayed on separate tabs.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click the  **New tab** icon beside the tab name.

A new tab is added. The tab is displayed as a separate dashboard tab in view mode.

5.4.2 Delete dashboard tabs

You can delete any view from a dashboard. A dashboard tab is displayed as a separate tab in the dashboard editor.

Warning

Deleted dashboard tabs cannot be restored.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click  **Show menu** beside the relevant tab title.
3. Click **Delete** in the pop-up menu.
4. Click **Delete**.

The selected dashboard tab is deleted.

5.4.3 Set dashboard tab properties

You can set the properties of a dashboard tab. You can specify a name for the dashboard tab, select a style and set the selected view as default.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a dashboard tab.
3. Click the  **Settings** icon beside the tab title.
4. Enter a text in the **Name** box in the pop-up menu.
5. To set the dashboard tab style, click the **Style** selection box and select a style.
This option is only available for smart dashboards.
6. To set the dashboard tab background, click the color selection box and select a background color.
This option is only available for fixed-grid dashboards.
7. Click **Set tab as default** to set the current selected tab as the default view.
The default view is displayed initially when you open the dashboard in view mode.

The dashboard tab properties are set.

5.4.4 Set dashboard tab style

You can change the style, for example, the background color, applied to the dashboard tab.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click  **Show menu** beside the tab title of the relevant tab.
3. Click the **Style** selection box and select a style.

The dashboard tab style is set.

5.4.5 Set dashboard tab as default

You can set the current selected tab as the default tab. The default tab is displayed initially when you open the dashboard in view mode.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click  **Show menu** beside the tab title of the relevant tab.
3. Click **Make default**.

The dashboard tab is set as default view.

5.5 Other

5.5.1 Use dynamic URL selection

You can select specific elements of widgets (for example, a grid row, a combobox entry, a pie slice, etc.) dynamically by specifying URL parameters. When you use the URL to open a dashboard in view mode, the elements are automatically preselected and, if applicable, the preselected elements are used as a filter or an action trigger.

You can select the URL parameters required for dynamic URL selection in the properties menu of a widget and copy them to the clipboard. You can set a data column preselection for each widget that supports data preselection, for example, Grid, Input field, or several charts.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click a widget that supports data preselection. The relevant properties dialog is displayed. Before you set a preselection you must assign the relevant data columns to the widget. See Assign data sources to widgets (Seite 31).
3. Click the **URL selection** tab.
4. You can enter a widget ID in the **URL-ID** input field. If you change the preset ID, you must choose a unique ID within your dashboard and you must save the dashboard to keep the new **URL-ID**.
5. The **URL example** field contains a URL selection of the widgets including all assigned data columns and corresponding dummy values.
6. You can edit the URL selection according to your requirements in the **URL example** field.

Expected formats:

Number: Do not use thousands separators and use dots as decimal separators.

Date / time: yyyy-MM-ddThh:mm:ss

Text: no restrictions (URL-encoded)

7. Copy the URL selection to the clipboard. If your browser supports this function, a **Copy to clipboard** button is available.
The URL parameters are copied to the clipboard of your operating system (URL-encoded).
8. Add the parameters to a URL that you want to use to open a dashboard and enter the URL in your web browser.
9. To open a specific dashboard tab, add the corresponding tab parameter to the dashboard URL.
 - a. Click  **Show menu** beside the tab title of the relevant tab.
 - b. Click the **URL selection** tab.
 - c. You can enter a tab ID in the **URL-ID** input field. If you change the preset ID, you must choose a unique ID within your dashboard.
 - d. Save your settings.
 - e. Copy the URL selection to the clipboard. If your browser supports this function, a **Copy to clipboard** button is available.
 - f. Add the URL selection that was copied to the clipboard to your dashboard URL.

The dashboard is displayed in view mode with the specified selection.

Example

Dashboard preselection parameters

```
&cn16.Time=2015-12-23&cn16.Location=New%20York
```

Tab preselection parameter

```
&tab=tab1
```

URL with parameters added

```
http://<local  
host>:8080/mashzone/hub/dashboard/dashboard.jsp?editmode=false&guid=0bd1cbcc-49d2-4  
cb1-a5fe-72cfdc624cda
```

```
&cn16.Time=2015-12-23&cn16.Location=New%20York&tab=tab1
```

These URL parameters are applied when you open the dashboard. To apply a modified preselection, you must reload the dashboard page. However, you can also apply a selection in an open dashboard without reloading the entire page. In this case, the selection string must begin with #... instead of &...

```
...#cn16.Time=2015-12-23&cn16.Location=New%20York
```

5.5.2 Hide tab bar in view mode

You can hide the tab bar in dashboard view mode.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Manage** in the dashboard main menu.

3. Click **Properties**.
4. Enable the **Hide tab bar in view mode** option.
5. Click **OK**.

Your settings are applied.

The tab bar is hidden when you open the dashboard in view mode. Only the active tab is displayed.

5.5.3 Set display size

You can set the display size for the current displayed dashboard. Set the zoom factor to improve the legibility of a dashboard.

The option is only available in fixed-grid dashboards. See [Switch to fixed-grid workspace](#) (page 15) for details.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Options** in the dashboard main menu.
 - Click **Zoom In** to increase the display size.
 - Click **Zoom Out** to decrease the display size.
 - Click **Set Zoom to 100%** to restore the default display size.

The display size of the current dashboard is adjusted accordingly.

5.5.4 Set grid lines

You can set the grid lines in the desktop background. The grid lines help you to arrange the individual widgets more easily. The widgets are automatically aligned to the grid lines.

These options are only available in fixed-grid dashboards. See [Switch to fixed-grid workspace](#) (page 15) for details.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Options** in the dashboard main menu.
3. Click **Increase Grid Space**.
4. Click **Reset Grid** to restore the default grid space.
5. Click **Decrease Grid Space**.

The grid lines of the current dashboard are adjusted accordingly.

5.5.5 Allow anonymous access to dashboards

MashZone NextGen allows the anonymous access to dashboards. A user can view a dashboard without logging in to MashZone NextGen.

Using the dashboard URL and the guest user account, you can allow a user to view a dashboard without logging in to MashZone NextGen. If a dashboard is opened with an anonymous access, only the **?** help button is available, and the user cannot switch to edit mode. The **Options** menu is only visible if the dashboard was created using the fixed grid layout (page 15).

Procedure

1. Add the guest account with view permission to the dashboard and the associated data feeds and aliases.

For details about managing dashboard permissions and assigning a user group to related data feeds and aliases, see Manage dashboard permissions (page 16).

- a. Open the **Manage dashboard permission** dialog for the dashboard.
- b. Add the **Presto_Guest** group to the Principals list.
The **View** permission is enabled for the guest user group by default.
- c. Enable the **Assign the relevant view permissions to related assets** option.
The **Presto_Guest** group is assigned to all associated data feeds and aliases.

2. Create the dashboard URL.

- a. Add the **x-p-anonymous=true** parameter to the dashboard URL in the address bar of your Web browser.

For example

```
http://localhost:8080/mashzone/hub/dashboard/dashboard.jsp?guid=f701601a-5d15&x-p-anonymous=true
```

- b. Copy the complete dashboard URL to the clipboard.

The required dashboard URL is available and can be used to view the dashboard anonymously.

5.5.5.1 PPM data source operator

If a dashboard or an associated data feed contains a PPM data source operator (page 309), a fixed basic authentication is required in the operator to allow anonymous access. Single Sign-On does not work with anonymous access.

5.5.5.2 File resources

If the dashboard contains an external resource file that was uploaded to MashZone NextGen, for example, an image, you must perform the following steps.

- Execute the following command in the MashZone NextGen API console (page 438) to grant the anonymous access.

```
{
  "version": "1.1",
  "sid": "PolicyService",
  "svcVersion": "0.1",
  "oid": "addPermissions",
  "params": [
    "[NAME_OF_THE_FILE]",
    "type.entity.file",
    "VIEW",
    [
      { "principalId": "Presto_Guest",
        "principalTypeId": "SpecialGroup"
      }
    ]
  ]
}
```

- Add the **x-p-anonymous=true** URL parameter to the source URL to access the uploaded file. For example, see the **Source URL** parameter of the Image (page 280) widget.

5.5.6 Customize the MashZone NextGen welcome page

You can replace the logo and the welcome text displayed on the MashZone NextGen welcome page. In addition, you have the option to change the appearance of the welcome page header.

Prerequisite

To edit the welcome text you must have administration privileges.

You must replace the relevant logo and welcome text on each hosting MashZone NextGen server.

Your changes in the welcome page header are applied to all application headers in MashZone NextGen, also to the dashboard and data feed editor.

Procedure

1. To replace the logo on the MashZone NextGen welcome page, you must replace the `landing_page_icon.png` graphic file in the following folder by your own graphic file.
<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\hub\assets\images\`</code>`
2. To replace the welcome text on the MashZone NextGen welcome page, open the `welcome-text.json` file with a text editor. Enter a new text and save your changes.

The file is located in the <MashZone NextGen installation>\apache-tomcat\webapps\mashzone\WEB-INF\config folder.

3. To change the appearance of the application header, you can edit the style templates supplied with MashZone NextGen. For details see Edit style templates (Seite 152).

4. Restart the MashZone NextGen server(s).

Your changes are applied.

If you have migrated the **application.less** file from MashZone NextGen version 9.12 or older and you have customized the appearance of the application header, you must adapt the **application.less** file of your current version.

You must change the path reference of the logo used in the application header. The path reference in the **@brand-logo** key must be absolute and starts with "/", for example, '/hub/dashboard/assets/images/my-logo.png'.

Add the following key.

```
// Font size used for application menu items in the dropdown menu of the masthead
@navigation-list-font-size-menu-dropdown: 14px;
```

5.5.7 Display dashboards without application header

In view mode, you can display your dashboards without the application header of MashZone NextGen. Only the dashboard is displayed in the web browser tab.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. You can hide the application header by adding the `appheader=false` parameter to the dashboard URL in the view mode in the following form:

http://<url to dashboard>?appheader=false

3. In addition, you can display a drop-down menu in the dashboard, which provides zoom, help and logout options. The  menu icon is displayed in the top right corner. Add the `showmenu=true` parameter to the dashboard URL in the following form.

http://<url to dashboard>?appheader=false&showmenu=true

The dashboard is displayed without the application header in view mode.

URL Example

http://<local host>:8080/mashzone/hub/dashboard/dashboard.jsp

?editmode=false&guid=0bd1cbcc-49d2-4cb1-a5fe-72cfdc624cda&appheader=false&showmenu=true

5.5.8 Export dashboards

You can export your MashZone NextGen dashboards.

You can use the export zip file to create a backup or to import your dashboards into another MashZone NextGen installation.

Procedure

1. Open a command window and move to the **MashZoneNG-install/prestocli/bin** folder.

2. Enter this command:

```
padmin exportDashboard -i identifier [-f output-file] [-l prestoURL] -u username -w password [-v] [-o]
```

- -i identifier: Mandatory dashboard identifier. It can be "id=", "name=" or "all", enclosed in quotes.
- -i "name=dashboardname": If there are multiple dashboards with the same name only the first dashboard found will be exported.
- -i "id=43243244434432": The dashboard ID (GUID) is unique in the MashZone NextGen system.
- -i "all": Exports all dashboards for that user.
- -f output-file: Optional path and name for the export. If omitted, an output zip file is created in the folder in which this command is executed:
Single export with option -i "id=3456" or "name=name" create a new file with name "name_guid.zip"
Multiple export with option -i "all" create a new file dashboard-export-timestamp.zip
- -l prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to **http://localhost:8080/mashzone/esd/api**.
- -u username: **MashZone NextGen user name to log in with. This account must have MashZone NextGen administrator permissions.**
- -w password: is the MashZone NextGen password to log in with.
- -v: is an optional flag to activate verbose logging.
- -o: Optional flag to overwrite an existing export file.

Once the export command completes successfully, you can use the output file to import dashboards into MashZone NextGen.

Permissions for each dashboard were automatically stored in the zip file. If no permissions are assigned to the dashboard, the permission file saved is empty.

The zip file also includes information about the dashboard creator.

5.5.9 Import dashboards

You can import dashboards in MashZone NextGen.

The dashboards are saved in a zip containing the dashboard definition, resource policy, and dashboard permissions, etc.. If you import a dashboard including permissions, the creator of the dashboard can view and edit the dashboard. The importer of a dashboard automatically becomes the creator of the dashboard if the dashboard is imported without permissions.

Procedure

1. Open a command window and move to the <MashZoneNG-installation>/prestocli/bin folder.
2. Enter this command:

```
padmin importDashboard [-l prestoURL] -f input-file -p importPermissions -u username -w password [-v] [-o]
```

- -f input-file: Path and name of the import zip file.
- -p importPermissions: Imports the resource policy and permissions saved in the import zip file.
- The importer of a dashboard automatically becomes the creator of the dashboard if the dashboard is imported without permissions. And only administrators can see and work with the dashboards imported.
- -o: is optional. Allows overwriting an existing dashboard in MashZone NextGen Dashboard.
- -l prestoUrl: Optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this results in **http://localhost:8080/mashzone/esd/api**.
- -u **username**: **is the MashZone NextGen user name to log in with. This account must** have MashZone NextGen administrator permissions.
- -w password: is the MashZone NextGen password to log in with.
- -v: is an optional flag to turn on verbose logging.

Once the import command completes successfully, you can use the imported dashboards in MashZone NextGen.

If you have imported dashboards from Presto 3.9 into MashZone NextGen, save the imported dashboards in edit mode of the widget before you display them in view mode. Otherwise, an error message is displayed.

5.5.10 Create widget style template

You can create your own widget style templates. Creating your own style templates enables you to customize the look and feel of your widgets, for example, colors schemes, fonts, or alignment. By default, a dashboard style template contains a set of widget styles, that is applied automatically, when you assign a template to a dashboard. By creating your own **less** template files you are able to overwrite the widget styles provided by the dashboard style template file. A widget template applies to all widgets of the same type, for example, line chart, of a dashboard assigned.

For each widget style, for example, number_left_aligned, a **less** style template file, for example, number_left_aligned.less, has to be created. The template files have to be saved in an specific widget subfolder of the relevant dashboard template folder.

```
..\assets\custom-look-and-feel\dashboard\<dashboard template name>\<widget type>
```

The following folder names can be used to create the relevant widget subfolders: **actionbutton**, **barchart**, **bubblechart**, **colorpalette**, **columnchart**, **container**, **grid**, **horizontalgauge**,

label, layoutgroup, layoutrow, linechart, piechart, speedometer, tab, and verticalgauge.

The dashboard style template files are located in the following folder on the MashZone NextGen server. The default dashboard template file is named **default.less**.

<MashZone NextGen

installation> \apache-tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\dashboard

Prerequisite

You have MashZone NextGen administrator permissions.

Procedure

1. Create a folder with the dashboard template name and a subfolder with the widget type name in the following form.

... \assets\custom-look-and-feel\dashboard\<dashboard template name>\<widget type>

2. Create a less style template file for each widget style you want to overwrite in the dashboard template.

You can add the styles you want to change or you can add new styles that are not part of the default style template.

3. Save your less template file in the widget type folder created.
4. Give the template file the name of the style you want to replace in the dashboard template.
5. Load the new style template file.
 - a. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
 - b. Click **Manage > Change style template** in the dashboard main menu.
 - c. Select the relevant dashboard style template in the dashboard drop-down menu and click **Update**.
6. Click OK.

The new widget style template is now available in the widget.

Example

You want to replace the **number_left_aligned** style of the **grid** widget type provided by the **default** dashboard template. Therefore, you have to create a folder for the default dashboard template and an addition subfolder for the grid widget:

... \assets\custom-look-and-feel\dashboard\default\grid

The content of the **number_left_aligned.less** template file might be as follows:

```
[presto-angulargrid] [ag-grid] { .number-cell{ text-align: left; } }
```

To assign the new style template to your widgets, see Change widget style template (page 70).

6 Use the data feed editor

The data feed editor as a graphical user interface is an easy-to-use tool to create, manage and view your data feeds without any programming knowledge.

To use the data feed editor, the **Dashboard administrator** function privilege is required.

To use data feeds on dashboards and in other data feeds, and to see the data on a dashboard, the **View** permission is required. See Manage data feed permissions (page 84) for details.

6.1 Create a data feed

You can create data feeds in the feed editor.

A data feed is a table that contains processed data. The data in the data feed table is calculated using a feed definition, which combines data from various data sources.

To create a data feed definition you must perform the following steps.

- Open a new data feed (page 81)
- Select data sources (page 81)
- Calculate the feed data (page 82)
- Add further operators (page 82)
- Connect the operators (page 83)
- Save the data feed (page 83)

6.1.1 Open a new data feed

You can create a new data feed by using the data feed editor.

Procedure

Click **Manage > New data feed** in the data feed editor main menu.

An new data feed is opened.

When creating a data feed, the **Output** element that finalizes the feed definition, is already set. The element is mandatory and cannot be deleted.

6.1.2 Select data sources

You can set one or more data sources for the data feed definition, for example, MS Excel, CSV, or XML files. The data sources for a data feed can be located locally, in the LAN, or on the Internet.

The data sources are represented by data source operators. Various options are available for setting the data source depending on the data source type.

See Data source operators (page 296) for a list of data sources and the relevant settings supported by the data feed editor.

The source files can be selected by specifying the path, and must be stored in the  Repository of ARIS Connect.

Procedure

1. Click the  symbol in the **Add data operations** bar if the symbol is not selected.
2. Click a data source, for example,  XML.
The selected operator is displayed in the data feed editor workspace.
3. Select the data source or connection type of the data source file required, for example, an URL or a local path to a XML file.
4. Enter a URL or a path to the data source file, for example, for an XML file. Or select a data source or a connection in a selection box, for example, a PPM connection.
5. Specify your additional settings.

The selected data source operator is inserted and the source data can be extracted.

6.1.3 Calculate the feed data

You can calculate the data for all operators of the feed definition and display the corresponding preview in the data feed table.

Procedure

1. Click the  **Calculate preview** icon in the header of an operator.
2. To read out the raw source data click the  **Calculate preview** icon in the data source operator.

The data feed is calculated up to the selected operator of the feed definition. The result is displayed in a preview table at the lower edge of the workspace.

6.1.4 Add further operators

Insert optional operators in the data feed definition to convert, calculate or transform data. For this the feed editor provides additional data transformation and user input operators.

You can use various operators to create calculation rules for calculating the data of your feeds.

See Data transformation operators (page 317) for a list of transformation operators and the relevant settings supported by the data feed editor.

Procedure

1. Click the  **Transformation** or  **Single values** icon in the **Add data operations** bar.
2. Click an operator or insert an operator using drag and drop.
The operator is displayed in the feed editor workspace.
3. Specify your settings.

The selected operator is inserted and configured.

6.1.5 Connect the operators

Connect the inserted operators to define the data flow of the data feed definition.

The data of an operator is forwarded to another operator using a link. The link is created as a connection between outgoing and incoming  anchor points of the individual operator. The permitted incoming anchor points are marked in blue for a selected outgoing anchor point.

User input operators are only connectable with **Single value** operators.

Procedure

1. Click the outgoing anchor point of a data source and drag it to an incoming anchor point of a transformation operator.
2. To disconnect two operators click an incoming or outgoing anchor point and drop it in the workspace.
3. Define the data flow among the inserted operators in a similar manner.
4. Connect the last transformation operator with the **Output** operator to finalize the data feed definition.

Your data feed definition is configured.

To view the calculation result of the completed data feed definition click the  **Calculate preview** symbol of the **Output** operator.

6.1.6 Save the data feed

You can save the data feed and give it a unique name.

Procedure

1. Click **Manage** > **Save** in the main menu.
You can create a copy of the current opened data feed with the **Save as** option.
2. Specify your settings.

The data feed is saved on the server.

You can change your settings by editing the data feed (page 83).

6.2 Edit a data feed

You can edit a data feed that is available in the data feed editor.

Procedure

1. Click **Manage** > **Open** in the data feed editor main menu.
2. Select an **Available data feed**.
3. Specify your settings.
4. Click **Manage** > **Save** in the data feed editor main menu.

Your settings are applied.

6.3 Delete data feeds

You can delete data feeds in the data feed editor.

Warning

Deleted data feeds cannot be restored.

Procedure

1. Create a dashboard (page 11) or open a dashboard (page 15) in the dashboard editor.
2. Click **Manage** > **Delete** in the data feed editor main menu.
3. Click **Delete**.

The selected data feed is deleted.

6.4 Edit data feed properties

You can edit the properties (name, description and tags) of existing data feeds.

Procedure

1. Create a new data feed (Seite 81) or open an existing data feed (page 83).
2. Click **Manage** > **Properties** in the feed editor main menu. The **Dashboard properties** dialog is displayed.
3. Enter the **Name** of the data feed.
4. Enter an optional **Description**.
5. Optionally, enter comma-separated search tags in the **Tags** field.
6. Click **OK**.

Your changes are applied.

6.5 Manage data feed permissions

You can manage the permissions of data feeds in the data feed editor. You can assign specific access permissions to individual users or user groups. If you assign permissions to a user group, the permissions are automatically assigned to all members of the user group.

For new users and user groups of a data feed, you can automatically assign view permissions to all associated assets of the data feed, such as aliases or other data feeds. It is not required to assign the permissions to each asset manually. A user requires the view permission for all associated assets to display the corresponding source data in the data feed. If view permissions are not assigned to all associated assets, a corresponding option to assign the missing view permissions is additionally displayed in the dialog.

You can specify the following access rights for saved data feeds.

- **Edit**

The user can use data feeds to create dashboards or to include them in other data feeds. The user can edit data feeds in the data feed editor.

- **View**

The user can use data feeds to create dashboards or to include them in other data feeds. The user can view the data of the data feed in view mode on the related dashboard.

Procedure

1. Create a new data feed (Seite 81) or open an existing data feed (page 83).
2. Click **Manage > Permissions** in the data feed editor main menu. The **Manage dashboard permissions** dialog is displayed.
3. Enter a term in the search field and click **Search**. If you click **Search** without specifying a search term, all users and user groups are listed.
4. Click **Show MashZone NextGen default groups** to show only default users or user group in the **Search results** field.
5. Drag a user or a user group from the **Search result** field and drop it on the **Principals with permissions** field.

By default, the creator of the data feed is shown in the **Principals with permissions** list.

6. Enable or disable the **View** or **Edit** privileges of a user or a user group.
7. Click **Save**.

The button is available if the option **Assign the relevant view permissions to related assets** is disabled, or the view permissions are already assigned to all associated assets.

8. Enable the option **Assign the relevant view permissions to related assets** to assign the required view permissions to all associated data feeds and aliases.

The option is available if the view permissions are not assigned to all associated assets.

9. Click **Next**.

A new dialog opens. The first list in the dialog contains assets whose view permissions you can update. The second list contains assets whose view permissions you cannot change. At least one of the following prerequisites must apply to change the view permissions for data feeds or aliases.

You are an administrator who can edit the permissions for aliases.

You have permissions to view and edit data feeds.

You have permissions to create and edit data feeds.

Your changes are applied.

If you want to remove a user or a user group from the Principals with permissions list, click the **Delete** icon. Deleted permissions for a data feed do not affect the associated data feeds or aliases.

6.6 Change the feed editor style

You can edit the style templates supplied with MashZone NextGen. Editing the style templates enables you to customize the look and feel of the feed editor, for example, colors schemes, fonts, brand logo or background colors.

Prerequisites

You have MashZone NextGen administrator privileges.

Procedure

1. Edit the style template file **application.less** located in the following folder on the MashZone NextGen server.
<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\application.
2. Reload the changed style template file.
 - a. Open the MashZone NextGen dashboard Editor.
 - b. Click **Manage > Change style template** in the main menu.
 - c. To reload the application style template click **Activate**.
 - d. Click **OK**.
3. Open the MashZone NextGen feed editor.

Your changes are applied.

6.7 Data feed operators

The data feed editor provides a wide range of data source and data transformation operators for creating data feed definitions (Seite 81).

- Data source operators (page 296)
- Data transformation operators (page 317)
- User input operators (page 357)

6.8 Use Materialized Feeds

Data feed processing usually caches results internally, rather than re-evaluating the data feed each time. The refresh of the cache depends on the configured refresh rates of the underlying data sources of the data feed. In some scenarios, you may need explicit control over the data feed results and you want to run follow-up analyzes over those specific results. For that purpose, you can materialize the results of a data feed evaluation in Terracotta DB, having explicit control over the storage point in time. Instead of re-evaluating the original data feed, you can work with the data feed results materialized in Terracotta DB.

Working with Materialized Feeds affects several components and workflows:

- You must configure a Terracotta DB connection that allows to create a dataset on the Terracotta DB server and materialize feed results in it.
For details on configuring a Terracotta DB connection, see [Manage Terracotta DB connections \(Seite 177\)](#).
- You must configure a Materialized Feed that particularly includes the definition of which data feed should be materialized in which dataset using which Terracotta DB connection.
For details on configuring a Materialized Feed, see [Manage Materialized Feeds \(page 188\)](#).
- You can either use the Admin Console, API console, the MashZone NextGen administration tool, or publicly available REST webservices in order to calculate data feed results and materialize them subsequently in Terracotta DB.
For details on triggering a Materialized Feed, see [Materialize data feed results in Terracotta DB \(page 191\)](#).
- You can access and analyze data feed results materialized in Terracotta DB using the Terracotta DB source operator ([page 311](#)).
- You can define which user or user group can materialize data feed results in a dataset, access those datasets during development, and access the data feed results in a dashboard.
For details on configuring permissions on a Materialized Feed, see [Share Materialized Feeds \(page 190\)](#).

6.8.1 Use datasets storing data feed results

To access a dataset with Materialized Feed results on a Terracotta DB server, you have to configure a Terracotta DB connection and a Materialized Feed in MashZone NextGen in advance. Both can be defined in the MashZone NextGen Admin console. See [Manage Terracotta DB connections \(Seite 177\)](#) and [Manage Materialized Feeds \(page 188\)](#).

Procedure

1. Create a new data feed ([Seite 81](#)) or open an existing data feed ([page 83](#)).
2. Insert a **Terracotta DB** ([page 311](#)) data source. For details, see [Select data sources \(page 81\)](#).
3. Select **Materialized Feed** in the **Source** drop-down menu.
4. Select a Materialized Feed alias in the **Materialized Feed** drop-down menu. The selected alias specifies the Materialized Feed with the dataset containing the Materialized Feed results. The menu contains all available Materialized Feed aliases for which you have the **Usage** privilege. The corresponding dataset is displayed in **Dataset Name** box.
5. Configure the RAQL query, refresh rate, and input parameters. For details, see **Terracotta DB** ([page 311](#)) data source.

Your settings are applied.

6.8.2 Administrate Materialized Feeds in Terracotta DB

Materialized Feeds are used to calculate feed results and materialize them in a dataset which is managed by a Terracotta DB server. MashZone NextGen creates those datasets if not already present. MashZone NextGen can also delete them if required. In a few use cases deletion of those datasets is not possible:

- A Materialized Feed is deleted (page 190) with the option to delete the corresponding dataset being activated. The dataset will not be deleted if it is still referenced by another Materialized Feed.
- When a Materialized Feed is updated (page 189) and the name of the dataset is changed, the old dataset will not be deleted.

If you want to delete these datasets, a Terracotta DB administrator must perform that task using Terracotta DB tooling.

Besides Materialized Feed datasets, MashZone NextGen also creates a dataset to maintain metadata associated with Materialized Feed datasets as well as runtime information. This metadata dataset must neither be manipulated nor deleted as it is required for the correct interaction of MashZone NextGen with Materialized Feed datasets.

6.8.3 Use datasets from external application

To access a dataset on a Terracotta DB server, which has been generated by an external application, you must configure a Terracotta DB connection in MashZone NextGen in advance. Terracotta DB connections can be defined in the MashZone NextGen Admin console. For details, see Manage Terracotta DB connections (Seite 177).

Procedure

1. Create a new data feed (Seite 81) or open an existing data feed (page 83).
2. Insert a **Terracotta DB** (page 311) data source. For details, see Select data sources (page 81).
3. Select **Dataset** in the **Source** drop-down menu.
4. Select a Terracotta DB connection in the **Terracotta DB alias** drop-down menu. The alias specifies the Terracotta DB server you want to connect with. The menu contains all available Terracotta DB connections you have **Usage** privilege for.
5. Select a dataset alias in the **Dataset name** drop-down menu. The menu contains all available Terracotta DB datasets for the selected Terracotta DB alias. For details, see Manage Terracotta DB connections (Seite 177). Datasets on the Terracotta DB server generated by MashZone NextGen for Materialized Feeds are excluded from the list.
6. Click **Configure columns** to configure the columns of the dataset.

When selecting a dataset, schema guessing is used to derive its columns. Using 1000 records, the guessing mechanism derives the set of all unique cell definitions of those records and defines corresponding columns.

- a. Unselect a column to exclude it from processing.
 - b. Select or enter a (positive) **Sample size** and click **Derive columns** to re-run the schema guessing with a larger sample size. The current configuration is omitted. The larger the sample size is, the longer the schema guessing takes.
Click **Add column** to manually add a new column.
 - c. Enter a unique column name and select a data type for the new column.
The combination of name and type must fit to the cell definition in the store. If not, the corresponding values will always be null. For example, if the cell definition is age with type INT and you specify age with type DOUBLE, the column definition does not fit to the cell definition and the values are null.
7. Configure the RAQL query, refresh rate, and input parameters. For details, see **Terracotta DB** (page 311) data source.
 8. Your settings are applied.

6.8.4 Which privileges are required?

In addition to configuring the sharing settings for Materialized Feeds (page 190), you must also configure the required privileges for the data feed and the Terracotta DB connection used.

When working with Materialized Feeds, the following use cases are relevant:

- In dashboards mode the contents of a dataset with Materialized Feed results.
 - Data feed accessing the dataset with Materialized Feed results: Enable the **Display** privilege (page 84).
 - Materialized Feed: Enable the **Display** privilege (page 190).
 - Terracotta DB connection: Enable the **Display** privilege (Seite 179).
- In feed editor provide the Materialized Feed alias in the list of available aliases in the Terracotta DB source operator.
 - Materialized Feed: Enable **Usage** privilege (page 190).
 - Terracotta DB connection: Enable **Usage** privilege (Seite 179).
- Calculate a data feed and materialize its results in a dataset being managed by a Terracotta DB connection. This can only be done by users with administration privilege, who per se have the following artifact rights.
 - Data feed whose results are materialized: Enable **Display** privilege (page 84).
 - Materialized Feed: Enable **Usage** privilege (page 190).
 - Terracotta DB connection: Enable **Usage** privilege (Seite 179).

6.9 Other

6.9.1 Import data feeds

You can import data feeds to MashZone NextGen.

The data feeds are saved in a zip file that contains among other things the data feed definition, resource policy and data feed permissions. If you import a data feed including the permissions then the creator of the data feed can view and edit the data feed. Importing data feeds without the relevant permissions makes the importer automatically to the creator of these data feeds.

Procedure

1. Open a command window and move to the MashZoneNG-install/prestocli/bin folder.
2. Enter this command:

```
padmin importFeed [-l prestoURL] -f input-file -p importPermissions -u username -w password [-v] [-o]
```

- -f input-file: path and name for the import zip file.
- -p importPermissions Imports the resource policy and permissions saved in the import zip file.
- If you import data feeds without permissions makes the importer automatically to the creator of these data feeds and the data feeds has no explicit permissions which means that only administrators can see and work with the data feeds .
- -o: is optional. Allows to overwrite an existing data feeds.
- -l prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to `http://localhost:8080/mashzone/esd/api`.
- -u username: is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.
- -w password: is the MashZone NextGen password to log in with.
- -v: is an optional flag to turn on verbose logging.

Once the import command completes successfully, you can use the imported data feeds in MashZone NextGen Feed Editor.

Example

```
pAdmin importFeed -l http://localhost:8080/mashzone/esd/api -f feedDefinition.zip -u Administrator -w manage -o
```

With this command the content of the data feed file "feedDefinition.zip " will be imported to MashZone NextGen.

6.9.2 Export data feeds

You can export your MashZone NextGen data feeds.

Export creates an export file that you can use to import data feeds to MashZone NextGen.

Procedure

1. Open a command window and move to the MashZoneNG-install/prestocli/bin folder.
2. Enter this command:

```
padmin exportFeed -i identifier [-f output-file] [-l prestoURL] -u username -w password [-v] [-o]
```

 - -i identifier: mandatory data feed identifier. It can be "id=", "name=" or "all", enclosed in quotes.
 - -i "name=feedname": If there are multiple data feeds with the same name then only the first founded data feed will be exported.
 - -i "id=43243244434432": The data feed id (Guid) is unique in the MashZone NextGen system.
 - -i "all": Export of all data feeds for that user.
 - -f output-file: an optional path and name for the export file to put data feeds. If omitted, this generates an output zip file in the folder where this command is executed:
 - Single export with option -i "id=3456" or "name=name" create a new file with name "name_guid.zip"
 - Multiple export with option -i "all" create a new file datafeed-export-timestamp.zip
 - This file must not already exist, unless you also use the **-o** option.
 - -l prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to `http://localhost:8080/mashzone/esd/api`.
 - -u username: is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.
 - -w password: is the MashZone NextGen password to log in with.
 - -v: is an optional flag to turn on verbose logging.
 - -o: an optional flag to overwrite an existing export file. If you omit this option, the output file must not already exist.

Once the export command completes successfully, you can use the output file to import data feeds to MashZone NextGen.

Permissions for each data feed were automatically stored into the zip file. If there are not any permissions assigned to the data feed an empty permission file is stored.

There is also an information about the data feed creator stored in the zip file.

Example

```
padmin exportFeed -l http://localhost:8080/mashzone/esd/api -f feedDefinition.zip -u Administrator -w manage -i "id=MyFeed"
```

This created zip file "feedDefinition.zip" contains all information of the data feed "MyFeed" incl. definition and permissions.

6.9.3 Display data feed without application header

You can display a data feed in the feed editor without the application header of MashZone NextGen.

Procedure

1. Create a new data feed (Seite 81) or open an existing data feed (page 83).
2. You can hide the application header by adding the `appheader=false` parameter to the data feed URL in the following form.

`http://<url to dashboard>?appheader=false`

3. In addition, you can display a drop-down menu in the data feed which provides zoom, help and logout options. The  menu icon is displayed in the top right corner. Add the **showmenu=true** parameter to the data feed URL in the following form.

`http://<url to dashboard>?appheader=false&showmenu=true`

The data feed is displayed without the application header in the feed editor.

URL example

`http://<local host>:8080/mashzone/hub/dashboard/feededitor.jsp?guid=32d64647-e3ba-403e-8406-5b289ec3d81d?appheader=false&showmenu=true`

7 Valuable information

This section provides background information to assist you in carrying out the relevant procedures.

7.1 What are dashboards?

Dashboards are interactive applications that collect, combine, and visualize data from different data sources, for example, ARIS table or CSV files. Dashboards are composed of individual widgets (for example, line chart or grid). They obtain their data from data sources and display it. Dashboards make it easy to visualize and analyze information. You can combine data from any original source and visualize them by means of graphic elements, filter the displayed results interactively and thus analyze them intuitively.

The dashboard editor as a graphical user interface gives you an easy graphic way to create, manage and view your dashboards. The dashboard editor provides the edit mode and the view mode. In the edit mode you can create and manage your dashboards. The view mode enables you to view and use your dashboards interactively.

See Use dashboards in view mode (page 5) for instructions.

See Create a dashboard (page 11) for instructions.

7.2 What are data feeds?

A data feed is a table containing prepared data. It consists of several columns that contain numerical values (for example, figures), text, or date values. Each row in the calculated result of a data feed corresponds to one data record.

The data in a data feed is calculated based on various data sources (for example, data from MS Excel, CSV, or XML files) by means of feed definitions. The source data is not an integral part of the data feed, but remains in its original sources, ensuring that it is constantly up-to-date. In addition to the external data sources, direct user entries in the data feeds can also be processed. Data feeds are used as data sources for dashboards.

Only one data feed can be assigned to each widget, with the same data feed being able to supply the data for multiple widgets. See Assign data sources to widgets (Seite 31) for details.

Feed definitions aggregate, extend, transform, or calculate data from one or more data sources. A feed definition can consist of any number of operators and data sources that are linked by connections. Data is calculated for each data source and each operator and then passed on to the operators linked to them for further processing. A feed definition delivers a data structure in the form of a list table as its result. All individual processing steps in the feed definition are based on this data structure.

The feed editor as a graphical user interface offers you an easy, visual way to create, manage, and view your data feeds, without programming knowledge being required. The feed editor

provides you with all supported data source operators, all relevant data transformation, and user input operators. The rule definitions can be created using drag and drop.

See Create data feeds (Seite 81) for instructions.

Data feed definition and feed table

Calculation result of operator 'Data feed'	
Customer	Revenue
SAP	10000.0
Siemens	20000.0
HP	15000.0
Volkswagen	12000.0

7.3 What is the view-only license?

Depending on the MashZone NextGen license used, various functionalities are available.

With MashZone NextGen version 10.3 you can use MashZone NextGen with the view-only license.

The view-only license allows the MashZone NextGen user to view only the contents of dashboards and data feeds.

- The user can view dashboards in view mode only.
- The user can view the data feed definition in the feed editor.
- In addition, the user can export and import dashboards and data feeds.
- If the user has **Edit permission** privileges for dashboards and/or data feeds, he can Manage dashboard permissions (page 16) and/or Manage data feed permissions (page 84).
- The user is not permitted to create and edit dashboards and data feeds.

You can use Command Central to manage licenses in MashZone NextGen (page 147).

8 Administration

8.1 Open the Admin Console

MashZone NextGen includes the Admin Console, a web based tool for MashZone NextGen administrators to configure and manage the MashZone NextGen Server. This simple, easy to use tool offers a wide degree of control to inspect, configure, and manage the features of MashZone NextGen.

Procedure

1. Click the  user icon in the program bar.
2. Click **Admin Console**.

The Admin Console opens.

8.2 MashZone NextGen Security

MashZone NextGen provides control of user interactions to register or create dashboards and data feeds, and secure access for all users to work with these artifacts based on policies that you define.

- **Change password:** For reasons of security, we strongly recommend that the MashZone NextGen administrator should change the standard MashZone NextGen password after installation. See Change technical user password (Seite 97).
- **Change password of target data sources:** For reasons of security, we strongly recommend to change the key that is used to encrypt or decrypt passwords of target data sources (for example, source operators, URL aliases, JDBC configurations). The key is included in the **authTokenKey** file located in <MashZone NextGen installation>/webapps/mashzone/WEB-INF/classes/. It can be changed by using the **padmin generateKey -a AES -f authTokenKey** command that creates a new **authTokenKey** file. First of all we recommend to create a backup of the existing **authTokenKey** file and then to copy the new file to that folder. The file should only be changed with an empty repository, as already encrypted passwords can not be decrypted any longer. The same applies to exported content. The system where the content should be imported, has to use the same key to be able to decrypt the passwords.
- **User Authentication:** based on the protocols shown above. You can also allow anonymous access if needed. See Authentication and Guest Access (page 98) for details.
- Incorporate password policies and expiring passwords.

Please consider the following security-relevant aspects :

- Always keep your operating system, installed widgets and applications updated. Run necessary security updates on a regular basis, in particular for your Web-Browser and installed plug-ins.

- Always keep your MashZone NextGen installation updated. Regularly check if new fixes are available for your installation and install them.
- To prevent unauthorized access to your system, only a limited number of users should be granted direct system access (for example, remote RDP access or directly via a management console).
- Limit network access by operating the server widgets behind a firewall. Only necessary services should be open in the firewall (for example, database).
- Hide network ports used solely for internal communication between server widgets.
- Set up a secured communication between client and server using HTTPS. For details, see [Configure HTTPS and Certificate Stores in the Application Server](#) (page 162).
- Install the latest security updates of your operating system, browsers and plug-ins, for example, Adobe Flash.

8.2.1 Manage your MashZone NextGen profile

Your user profile shows basic information about your account in MashZone NextGen and allows you to:

- Manage your locale and account information (page 96), if permitted
- Change your password (page 96), if permitted

8.2.1.1 Manage your locale and account information

In most cases, you cannot update any other account information because this comes from account information for your entire organization. In development or test environments where your account information is stored in the default MashZone NextGen Repository, you can save changes to this information.

Procedure

1. In the program bar, click the user name by which you are logged in to MashZone NextGen.
2. Click **About Me**.
3. Make your settings.
4. Click **Save changes**.

Your settings are applied.

8.2.1.2 Change your password

In most cases, you **cannot** update your password in MashZone NextGen because this comes from account information for your entire organization. In development or test environments where your account information is stored in the default MashZone NextGen Repository, however, you can reset your password from your profile.

Procedure

1. In the program bar, click the user name by which you are logged in to MashZone NextGen.
2. Click **My Password**.
3. Enter your new password and confirm this.
4. Click **Update Password**.

Your password is updated.

8.2.2 MashZone NextGen server

Security: this includes both authentication and authorization for users when dashboards and data feeds are viewed or run. The MashZone NextGen server also handles authentication with dashboards and data feeds information sources when they are run.

The MashZone NextGen server is integrated with your user repository (page 97) or identity server for user authentication. This can be basic authentication, secure connections and certificates or a single sign-on solution.

You define authorization policies for MashZone NextGen resources determining who can view or run dashboards and data feeds. Generally, users must be authenticated, but you can also define unlimited access, allowing 'guest' users without authentication to work with apps that are published to web sites, wikis or other environments.

8.2.3 MashZone NextGen repository

The MashZone NextGen repository contains information on users and groups, authorization policies, server configuration, notifications and much more.

User Data: for authentication and determining authorization.

Typically user data comes from your organization's LDAP directory which you integrate with MashZone NextGen. This may also use a single sign-on solution and an identity manager. However, MashZone NextGen also has a built-in user repository which you may use. User or group meta-data from LDAP allows MashZone NextGen to relate authorization policies with users.

8.2.4 Change technical user password

For reasons of security we strongly recommend that the MashZone NextGen administrator should change the standard technical user password after installation. The technical user password is encrypt and stored in two modules. You have to change both occurrences.

```
<MashZoneNG-install>/apache-tomcat/webapps/mashzone/WEB-INF/classes/mz.properties
```

```
<MashZoneNG-install>/apache-tomcat/webapps/mashzone/WEB-INF/mashzone.properties
```

Note: This procedure is only required for MashZone NextGen 3.9.01.

Procedure

1. Change the password in `.../mz.properties`.
 - a. Encrypt a new password using the `padmin` tool. Open the command line and enter following command. Replace the variable `<password>` by your new password, for example, `newPassword`.
\$ `<MashZoneNG-install>/prestocli/bin/padmin encryptProperty -u Administrator -w manage -p <password>`
 - b. Copy the output of the command line into `mz.properties`, for example,
`{ENC}A+yyI2FYYBy33IgNCWGQIQ==`.
`mzServer.secrete={ENC}A+yyI2FYYBy33IgNCWGQIQ==`
2. Change the password in `.../mashzone.properties`.
 - a. Encrypt a new password using the `encryptpassword` tool. Open the command line and enter following command. Replace the variable `<password>` by your new password, for example, `newPassword`.
\$ `<MashZoneNG-install>mashzone/tools/runtool encryptpassword -password <password>`
 - b. Copy the output of the command line into `mashzone.properties`, for example,
`46f712a61dc8d7ed244bf0ffd266ae1e`.
`presto.basicAuthPassword=46f712a61dc8d7ed244bf0ffd266ae1e`

The technical user password is changed.

8.2.5 Authentication and Guest Access

MashZone NextGen accepts requests from **both** unauthenticated (guests) and authenticated users.

Authentication is required:

- To use any feature in any MashZone NextGen Add-On that accesses the MashZone NextGen Server, unless that Add-On also supports guest access.

Requests are rejected with an authentication error when they do not provide one of:

- A valid MashZone NextGen session cookie. Sessions that have timed out are rejected with an appropriate error. See [Sessions and Timeouts](#) (page 99) for more information.
- Valid credentials. See [Valid Credentials](#) (page 99) for more information.
- Guest access header or parameter information.

8.2.5.1 User Authentication

MashZone NextGen is initially installed with a set of Default User Accounts (page 100) that you can use to get started. You configure MashZone NextGen to work with your LDAP Directory or you can continue to use the Default User Repository and simply add users and user groups to

MashZone NextGen. See [Use the Default MashZone NextGen User Repository](#) (page 132), [Manage Users](#) (page 133) and [Manage User Groups](#) (page 134) for more information.

Authentication to verify user identities is performed against LDAP or the default User Repository and uses one of these protocols:

- Basic authentication with username and password
- This is the default authentication mechanism. No additional configuration is needed.
- SSL and User Certificates
- A configurable Single Sign-On solution
- See [Authentication with Single Sign-On Solutions](#) (page 100) for configuration instructions.

Permission to work with MashZone NextGen artifacts can also be granted to guests (unauthenticated users), if needed.

8.2.5.2 Valid Credentials

When authentication is required, requests must have a valid MashZone NextGen session for an existing authenticated user or must supply either user credentials or digital certificate for authentication or an SSO token or ticket for a user that has been authenticated by the SSO solution. MashZone NextGen uses certificates, tokens or tickets to obtain the user's identity.

MashZone NextGen supports the following mechanisms to obtain user credentials or user IDs:

- Basic authentication using username and passwords. This is authenticated against the MashZone NextGen User Repository which may be a database or your LDAP Directory. See [Use the Default MashZone NextGen User Repository](#) (page 132) for more information.
- SSL and Certificate authentication where the user identifier in certificate information is configurable. This is authenticated against the MashZone NextGen User Repository which may be a database or your LDAP Directory, **unless** Dynamic User Support is enabled. See [Use the Default MashZone NextGen User Repository](#) (page 132) for more information.
- Single Sign-On (SSO) solutions which are configurable. With SSO enabled, MashZone NextGen delegates authentication to the SSO solution. Typically, configuration identifies an SSO token, ticket or cookie that MashZone NextGen uses to verify authentication with the SSO solution and to obtain the user ID. See [Authentication with Single Sign-On Solutions](#) (page 100) for more information.

If an authenticated request has no MashZone NextGen session, MashZone NextGen starts a new session and generates a MashZone NextGen session cookie. See [Sessions and Timeouts](#) (page 99) for more information.

8.2.5.3 Sessions and Timeouts

MashZone NextGen is based on the standard J2EE session mechanism supported by your application server. MashZone NextGen maintains a separate HTTP session for each authenticated

user that has a unique session cookie. Each request with a valid MashZone NextGen session cookie extends the timeout limit for that user session.

SSO solutions maintain their own sessions and may use their own session cookies. SSO session cookies can be used to authenticate users in MashZone NextGen. SSO sessions and MashZone NextGen session are separate.

The default session timeout for MashZone NextGen is 30 minutes, defined in the web-apps-home/mashzone/WEB-INF/web.xml configuration file. In general, HTTP session timeouts can be configured in **web.xml**, unless the application server provides other configuration mechanisms. Please see your application server documentation for additional information on session configuration.

8.2.6 Default User Accounts

MashZone NextGen has four user accounts that you can use 'out-of-the-box' to access MashZone NextGen dashboards and data feeds. These default users also illustrate the basic permissions to features in MashZone NextGen. See Built-In MashZone NextGen User Groups and Permissions (page 105) for more information on permissions.

Username	Password	Built-in Group / Permissions	Description
Administrator	manage	Presto_Administrator	A MashZone NextGen administrator.
dev	devdev	Presto_Developer	A developer.
power	powerpower	Presto_PowerUser	A domain expert or power user.
user	useruser	Presto_AuthenticatedUser	An end user or any user in the MashZone NextGen User Repository.

If you configure MashZone NextGen to use your LDAP Directory as the MashZone NextGen User Repository, these default user accounts are automatically disabled. If you use the Default User Repository, you can delete these user accounts in the Admin Console.

Important: You must make sure that at least one user account has MashZone NextGen administrator permissions.

8.2.7 Authentication with Single Sign-On Solutions

With a single sign-on (SSO) solution, MashZone NextGen delegates authentication to the SSO layer. MashZone NextGen has the following pre-configured options to integrate with SSO solutions:

- Agent-based SSO solutions, such as Netegrity SiteMinder. See Configuration for Agent-Based SSO Solutions (page 101) for instructions.

- MashZone NextGen provides the integration under My webMethods in a SSO scenario by SAML (Security Assertion Markup Language).
- See SSO integration in My webMethods (page 103) for details.

8.2.7.1 Configuration for Agent-Based SSO Solutions

MashZone NextGen delegates authentication to the SSO layer, but expects user identity information from the SSO layer in the request in either an HTTP header or a parameter in the request URL. MashZone NextGen uses an extractor to find identity information in the header or parameter, and uses a transformer, to derive the user ID from the identity information.

MashZone NextGen then uses the user ID to perform authorization and process the request.

To configure MashZone NextGen to work with an agent-based SSO layer, you configure the extractor and the transformer layers to work with your SSO solution and the identity information for your environment. MashZone NextGen provides a default extractor that looks for an HTTP header or parameter by name. MashZone NextGen also provides default transformers that handles cases where the identity information is just the user ID or can be found within the identity information using a regular expression.

You can also implement custom extraction or transformation layers to integrate MashZone NextGen with your SSO solution. See [Implementing a Custom SSO Filter](#) (page 103) for details.

Procedure

1. If needed, configure the MashZone NextGen User Repository. See [Use the Default MashZone NextGen User Repository](#) (page 132) for more information.

In previous releases, MashZone NextGen only supported SSO solutions with LDAP as the MashZone NextGen User Repository. This restriction no longer applies.

2. Change the SSO filter in the `applicationContext-security.xml` configuration file for the MashZone NextGen Server:

- a. Open `applicationContext-security.xml` in any text or XML editor.

This file is located in the `web-apps-home/mashzone/WEB-INF/classes` folder.

- b. Comment out the SSO filter bean (`<bean id="ssoProcessingFilter">`) for agent-based solutions (`class="com.jackbe.jbp.sas.security.ui.sso.SSONullPreAuthenticatedFilter"`).

For example:

```
<bean id="ssoProcessingFilter" > <property name="authenticationManager"
ref="authenticationManager" /> <property
name="continueFilterChainOnUnsuccessfulAuthentication" value="true" /> ...
</bean>
```

Comment out the complete XML element with its children.

- c. Comment in the bean `<bean id="ssoProcessingFilter" class="com.jackbe.jbp.sas.security.ui.sso.SSOPreAuthenticatedFilter">`.

Comment in the complete XML element with its children.

3. In the agent-based SSO filter bean, configure the `principalExtractor` property:

The default extractor uses a bean with the **HttpHeaderOrParamTokenExtractor** class.

```
<bean id="ssoProcessingFilter" > <property name="authenticationManager"
ref="authenticationManager" /> <property
name="continueFilterChainOnUnsuccessfulAuthentication" value="true" /> <property
name="principalExtractor"> <bean > <property name="httpHeaderName"
value="SM_USER"/> </bean> </property> ... </bean>
```

Change the value of the **httpHeaderName** property for this extractor bean to the name of the HTTP header or parameter that contains user identify information from your SSO solution.

If you have a custom extractor class, replace the default extractor bean with configuration for your custom class.

4. In the agent-based SSO filter bean, configure the **principalTransformer** property:

The default transformer property uses a bean with the

RegexExtractionStringTransformation class. This uses a regular expression to extract some portion of the value for the SSO header or parameter to get the final user ID that MashZone NextGen can use for authorization checks.

```
<bean id="ssoProcessingFilter" > <property name="authenticationManager"
ref="authenticationManager" /> <property
name="continueFilterChainOnUnsuccessfulAuthentication" value="true" />
<property name="principalExtractor"> <bean > <property name="httpHeaderName"
value="SM_USER"/> </bean> </property> <property name="principalTransformation">
<bean > <constructor-arg index="0" value="CN=(.*?),"/> </bean> </property>
</bean>
```

If the value of the SSO solution header or parameter contains more than just the user ID, for example a full DN from LDAP for a user, you can change the regular expression in the **<constructor-arg/>** parameter for the default bean to extract the user ID. The default regular expression extracts the CN portion of a user DN from an LDAP Directory.

If the value of the SSO solution header or parameter is **just** the user ID, no further transformation is needed. Change the **principalTransformer** bean to do nothing using the **NoOpStringTransformation** bean:

```
<bean id="ssoProcessingFilter" > <property name="authenticationManager"
ref="authenticationManager" /> <property
name="continueFilterChainOnUnsuccessfulAuthentication" value="true" />
<property name="principalExtractor"> <bean > <property name="httpHeaderName"
value="SM_USER"/> </bean> </property> <property name="principalTransformation">
<bean /> </property> </bean>
```

If you have a custom transformation class, replace the default transformer bean with configuration for your custom class.

5. Save this file and restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.2.7.2 Implementing a Custom SSO Filter

If the default extractor and transformer filters available in MashZone NextGen do not provide the functionality needed to allow MashZone NextGen to work with your SSO solution, you can create custom filters using the MashZone NextGen SSO Filter API MashZone NextGen SSO Filter API.

Procedure

1. Add the following JARs and classes to your classpath:

Classes in the web-apps-home/mashzone/WEB-INF/classes folder.

The web-apps-home/mashzone/WEB-INF/lib/presto_common.jar file.

2. Implement one or both filters:

To create a custom extractor, implement the **SSOTokenExtractor** interface, typically using the **AbstractSSOTokenExtractor** base class.

To create a custom transformer, implement the **Transformation** interface.

3. Add these classes to the classpath. Copy either the compiled class file or a JAR containing the compiled class file to one of these folders, respectively:

The external configuration folder, if any, for the MashZone NextGen Server. See Setting Up an External MashZone NextGen Configuration Folder (page 213) for more information.

Important: Deploying additional resources, such as custom SSO filters, to an external configuration folder simplifies future deployments or MashZone NextGen Server clusters.

web-apps-home/mashzone/WEB-INF/classes. This is the default location, but is not recommended as it complicates MashZone NextGen Server deployments.

web-apps-home/mashzone/WEB-INF/lib. This is the default location, but is not recommended as it complicates MashZone NextGen Server deployments.

8.2.7.3 SSO integration in My webMethods

You can integrate MashZone NextGen under My webMethods in an SSO scenario by SAML (Security Assertion Markup Language).

MashZone NextGen can accept SAML tokens for authentication in a SSO environment.

Specifically, My webMethods can act as an Identity Provider (IdP).

MashZone NextGen verifies the signature used to sign the SAML assertion is trusted by looking the comparing the signature to the **platform_truststore.jks** file. This file is a Java Keystore file, and can be managed using the Java "keytool" command. If the certificate used to sign the SAML assertion is not present in the **platform_truststore.jks** file, the assertion is rejected. The **platform_truststore.jks** file is configurable in

SAG_HOME/MashZoneNG/apache-tomcat/webapps/mashzone/WEB-INF/classes/presto.config.

Information on the Java "keytool" command can be found in the Java documentation:

<http://docs.oracle.com/javase/8/docs/technotes/tools/windows/keytool.html>.

Procedure

1. Within the presto.config file, the saml.truststore.file parameter contains the full path to the file. The default configuration uses the SAG_HOME/common/conf/platform_truststore.jks file. By default, the file contains the certificate used to sign My webMethods SAML assertions. No further configuration is needed in the My webMethods SAML case.
2. Within the presto.config file, the saml.truststore.passwd parameter contains the keystore password. The default configuration uses the password for the SAG_HOME/common/conf/platform_truststore.jks file. The default password is manage.
3. To accept SAML assertions signed by a third party, the signing certificate must be either imported as a "trusted certificate" to the currently configured platform_truststore.jks file, or the presto.config file must be altered to point to a different keystore file, where this signing certificate is already imported as a "trusted certificate".

8.2.8 Authorization Policies and Permissions

Authorization policies determine the actions that users can perform.

By default, authorization is enabled in MashZone NextGen. All actions are forbidden unless explicitly granted in a policy.

The categories of authorization policies that are defined in MashZone NextGen are shown below.

- **Access/Create Permissions:** are defined using MashZone NextGen built-in user groups as the principals. See the Built-In MashZone NextGen User Groups and Permissions (page 105) topic for detailed information these policies.
- To grant access to MashZone NextGen tools and enable users to create artifacts, you add users to these built-in groups. See Grant User Access to MashZone NextGen with Built-in Groups (page 104) for instructions.
- **Owner/Admin Permissions:** users automatically obtain owner permissions when they create artifacts. Administrator permissions are defined when you assign users to the **Presto_Administrator** built-in group (see Access/Create policies).
- Owners have full permissions to all actions for the artifacts they create. Administrators have owner permissions for **all** artifacts.
- You can also grant guest access to use artifacts. Guest access grants permission for anyone to run that artifact, even users who are not logged in. See Authentication and Guest Access (page 98) for instructions.

8.2.8.1 Grant User Access to MashZone NextGen with Built-in Groups

All users in the MashZone NextGen User Repository automatically belong to the **Presto_AuthenticatedUsers** built-in group which has permission to access the MashZone

NextGen welcome page (page 3). To enable users to work in MashZone NextGen, you must add them to the **Presto_PowerUser**, **Presto_Developer** or **Presto_Administrator** groups.

See the Built-In MashZone NextGen User Groups and Permissions (page 105) topic for information on the specific access policies for these groups. Or use the Default User Accounts (page 100) in MashZone NextGen to better understand the permissions for these groups, or give them the respective access rights.

Procedure

1. If you are using the Default User Repository with MashZone NextGen, both groups and users are defined with the Admin Console. To grant users permissions with the MashZone NextGen built-in groups:
 - a. Add users to the MashZone NextGen Repository. See Create Users (page 133) for instructions.
 - b. Assign users to the appropriate built-in groups. See Edit, Grant Permissions and other User Management Tasks (page 134) for instructions.
 - c. If desired, you can also automatically add users as members to groups when you create users. See Automatically Assign New Users to Groups (page 135) for instructions.
2. If you have configured MashZone NextGen to use your LDAP Directory as the User Repository, you relate users to the MashZone NextGen built-in groups in LDAP. To grant users permissions with the MashZone NextGen built-in groups:
 - a. Add Presto_Administrator, Presto_Developer and Presto_PowerUser as new groups in LDAP.

Note: To map users and groups in LDAP to MashZone NextGen built-in permissions, you add these predefined names to your LDAP Directory. Mapping from configuration in MashZone NextGen based on LDAP attributes is possible. Or defining alias names for these built-in groups is also possible. For more information and assistance, please contact your Software AG sales representative.
 - b. Assign users to these new groups in LDAP.

8.2.9 Built-In MashZone NextGen User Groups and Permissions

MashZone NextGen has a set of built-in user groups that define access permissions to the various features in MashZone NextGen. These built-in groups also define permissions for all artifact actions.

For more details on the permissions for these built-in groups, see Access Policies Using MashZone NextGen Built-In Groups (page 106).

8.2.9.1 Access Policies Using MashZone NextGen Built-In Groups

- **Guests** = users who are not authenticated. Guests can work with dashboards if the dashboards and all other artifacts that it depends on have granted permissions to the **Presto_Guest** built-in group.
- The most common use is to allow dashboards to be shown in public web sites or other environments where secure access is not needed.
Note: Granting guest access to dashboards does not implicitly grant permissions to the artifact to any authenticated MashZone NextGen.
- **End Users** = all authenticated users (in the MashZone NextGen Repository) that are not in another built-in group. Authenticated users can access MashZone NextGen welcome page to find artifacts, but they can only use the artifacts to which they have been granted run permissions. They also have **no access** to tools that create artifacts.
- **Developers** = users in the **Presto_Developer** group can find and create dashboards and data feeds using visual tools. Developers also have access to other technical information, such as the API Console.
- This group is typically used for IT or line-of-business developers involved in developing dashboards and data feeds for specific projects.
- **Administrators** = users in the **Presto_Administrator** group have unrestricted permissions in MashZone NextGen. They can work with any tools, features or artifacts. They also have permissions to use the Admin Console to configure and manage MashZone NextGen.

Administrators are the only built-in group that is required. You can use the other built-in groups to grant access to specific MashZone NextGen tools and features.

8.2.10 Protect MashZone NextGen Event Service access

You can use your own keystore and truststore to protect MashZone NextGen Event Service of unauthorized access.

After the installation, MashZone NextGen uses a default keystore and truststore. For security reason we recommend to change that configuration for production environments. Please make sure that the truststore, referenced by Event Service, contains the appropriate certificate for the key, referenced by MashZone NextGen. Event Service is only available if the configuration is valid.

If multiple MashZone NextGen nodes are used in a clustered scenario, it is recommended to use the same key for all MashZone NextGen instances.

The default keystore and truststore are located in the **common** and **conf** folders of the MashZone NextGen installation.

For authentication MashZone NextGen webapp sends an HTTP header "Authorization" with "Bearer [JWT]" as value.

Procedure

1. Edit the following parameters, used by MashZone NextGen, to use your own keystore file to generate the JWT required for authentication.

jwt.keystore.file

jwt.keystore.passwd

jwt.keystore.alias

The parameters are contained in the **presto.config** file in the following directory.

<MashzoneNG_install>\apache-tomcat\webapps\mashzone\WEB-INF\classes\

2. Edit the following parameters, used by Event Service, to use your own truststore file to verify the JWT.

rtbs.truststore.file

rtbs.truststore.passwd

The parameters are contained in the **rtbs.properties** file in the following directory.

<MashzoneNG_install>\rtbs\conf\

8.2.11 Anti-Clickjacking prevention when using iFrame

For security reason we recommend to configure your iFrame setting to protect your MashZone NextGen installation against clickjacking attacks.

Clickjacking is a vulnerability where an attacker creates a page that uses iFrame to render another page, then creates invisible controls on top of the rendered page that may be able to sniff user input.

General information on the clickjacking attack vector can be found on <https://www.owasp.org/index.php/Clickjacking>.

MashZone NextGen offers two ways to prevent successful clickjacking attacks. In order to allow iFrame on trusted sites, MashZone NextGen uses X-Frame-Options providing the **ALLOW-FROM** value. Using this, a website A can configure the header to carry the top level URI of a website B which is allowed to iframe website A. A second way to prevent clickjacking attacks is using the Content-Security-Policy that is supported by most web browsers.

Details on how to use iFrame with MashZone NextGen can be found in Embedding MashZone NextGen in external system environments (page 155).

8.2.11.1 MashZone NextGen HTTP header security filter

MashZone NextGen provides a specific HTTP header security filter included in the **web.xml** file. By default, this filter always sends the X-Frame-Option: **SAMEORIGIN**, that can be configured to send **ALLOW-FROM** to any number of trusted websites. This HTTP response header instructs the browser to refuse to render any content from MashZone NextGen in an iFrame, unless the iFrame is within MashZone NextGen itself.

HTTPHeaderSecurityFilter

Following the commented configuration in the **web.xml** file.

```
<filter> <filter-name>HTTP Header Security Filter</filter-name> <filter-class>
com.jackbe.jbp.sas.security.ui.http.HttpHeaderSecurityFilter </filter-class> <!--
Init Param: antiClickJackingEnabled Should the anti click-jacking header
(X-Frame-Options) be set on the response. Valid options: true or false When true,
X-Frame-Options will always contain "SAMEORIGIN". This instructs browsers to disallow
iframing of MzNG content outside of the MzNG application itself. If false,
X-Frame-Options will not be sent at all, which completely disables clickjacking
protection allows any site to iframe MzNG) Note: X-Frame-Options is superseded by
Content-Security-Policy.
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/X-Frame-Options -->
<init-param> <param-name>antiClickJackingEnabled</param-name>
<param-value>true</param-value> </init-param> <!-- Init Param: antiClickJackingUris
List of comma separated Uris for sites allowed to iframe content in MzNG. To allow
external sites to iframe MzNG content, uncomment this init param, and add the site
uri to the list. Also configure the 'Content Security Policy' filter below. If the
request to MzNG contains a referer value matching the scheme, host and port of one
of the Uris in the list, the X-Frame-Options header will send "ALLOW-FROM uri". This
allows the browser to render the iframe. If there is no match (or the list is empty)
X-Frame-Options will send "SAMEORIGIN" and the browser will refuse to render the iframe
Any site added to this list should also be added to 'Content Security Policy' header.
<init-param> <param-name>antiClickJackingUris</param-name>
<param-value>http://some-server.com</param-value> </init-param> --> <!-- Init
param: hstsEnabled Enable HTTP Strict Transport Security (HSTS) header
(Strict-Transport-Security) to be set on the response for secure requests -->
<init-param> <param-name>hstsEnabled</param-name> <param-value>true</param-value>
</init-param> <!-- Init Param: hstsMaxAgeSeconds The max age value that should be
used in the HSTS header. Negative values will be treated as zero. If not specified,
the default value of 0 will be used. --> <init-param>
<param-name>hstsMaxAgeSeconds</param-name> <param-value>604800</param-value>
</init-param> </filter> <filter-mapping> <filter-name>HTTP Header Security
Filter</filter-name> <url-pattern>/*</url-pattern> </filter-mapping>
```

The **antiClickJackingUris** parameters can take a list of comma separated URIs. The parameter is commented out by default. Any request for a MashZone NextGen resource containing a "Referer" header field matching the scheme, host and port of a URI in the **antiClickJackingUris** parameter will result in a response containing the X-Frame-Options response header with the appropriate **ALLOW-FROM** value. If there is no match, then the X-Frame-Options will carry the **SAMEORIGIN** value.

8.2.11.1.1 Example

The website <http://website-a.com> is configured as trusted, and therefore it is listed in the **antiClickJackingUris** parameter, and contains a page that uses iFrame to embed a MashZone NextGen dashboard. When a user visits this page on website-a.com, the browser will attempt to fetch the iFramed dashboard from MashZone NextGen. The request generated by the browser will carry the HTTP request header "Referer" containing the full URI to the page containing the iFrame. MashZone NextGen will match the "Referer" URI with the trusted URI from **antiClickJackingUris** parameter, and recognize that the website is trusted. As a result, the response will carry the HTTP response header "X-Frame-Options: ALLOW-FROM <http://website-a.com>". The browser will then allow the iFrame to render.

8.2.11.2 MashZone NextGen Content Security Policy

Most modern browsers such as Microsoft Edge, Chrome, Firefox and Safari check for the newer Content-Security-Policy HTTP header instead of X-Frame-Options. Within the MashZone NextGen **web.xml** file is a second HTTP filter class that sends the HTTP Header **Content-Security-Policy**. This filter is configured by default to send the value **frame-ancestors 'self'** which is equivalent to **SAMEORIGIN** in that it instructs the browser to only allow iFrame if the iFrame is already in the originating website.

The Content-Security-Policy is not supported by Microsoft Internet Explorer.

ContentSecurityPolicy

```
<filter> <!-- Allows setting of HTTP header Content-Security-Policy
http://www.w3.org/TR/CSP2/ To prevent clickjacking attacks default is
"frame-ancestors 'self'" which disallows external iframing of MzNG content. To allow
additional websites to iframe MzNG content, add the site Uri after 'self'. For example:
"frame-ancestors http://*.example.com/ 'self'"
https://developer.mozilla.org/en-US/docs/Web/HTTP/Headers/
Content-Security-Policy/frame-ancestors --> <filter-name>Content Security
Policy</filter-name>
<filter-class>com.jackbe.jbp.sas.security.ui.http.ContentSecurityPolicyFilter</
filter-class> <init-param> <param-name>policy</param-name>
<param-value>frame-ancestors 'self'</param-value> </init-param> </filter>
<filter-mapping> <filter-name>Content Security Policy</filter-name>
<url-pattern>/*</url-pattern> </filter-mapping>
```

8.2.11.3 Add a trusted site to allow iFrame

The default settings do not allow external sites to iframe internal MashZone NextGen assets such as dashboards, apps, etc. Specifically, "X-Frame-Options: SAMEORIGIN" and "Content-Security-Policy: frame-ancestors 'self'" are set, which instructs the browser to disallow rendering MashZone NextGen content in any external iFrame. Via configuration and re-start, we can relax this restriction.

Procedure

1. Open the **web.xml** file in a text editor. The file is located in <MashZone NextGen installation>/MashZoneNG/apache-tomcat/webapps/[presto|mashzone]/WEB-INF/.
2. Find the **<filter>** entry of the HTTP Header Security Filter and uncomment the **antiClickJackingUris** parameter.
3. Replace the sample URI ' http://some-server' with the URI of the website allowed to iframe MashZone NextGen content.
4. Find the **<filter>** entry for Content-Security-Policy. Insert the URI of the website allowed to iframe MashZone NextGen content into the **policy** parameter, between **frame-ancestors** and **'self'**

```
Example: <init-param> <param-name>policy</param-name>
<param-value>frame-ancestors http://*.eur.ad.sag:* 'self'</param-value>
</init-param>
```

8.2.11.4 Add multiple trusted sites to allow iFrame

To allow more than one website, perform the steps as shown in **Adding a trusted site to allow iFrame**.

Procedure

1. In the **HTTP Header Security** filter, add a comma separated list of URIs as the **antiClickJackingUris** value:

```
<init-param> <param-name>antiClickJackingUris</param-name>
<param-value>http://website-a.com, http://website-b.com:9999 </param-value>
</init-param>
```

2. In the **Content-Security-Policy** filter, add the URI to the policy parameter value, separated by a space:

```
<init-param> <param-name>policy</param-name> <param-value>frame-ancestors
http://website-a.com http://website-b.com 'self' </param-value> </init-param>
```

8.2.11.5 Content-Security-Policy using wildcards

The Content-Security-Policy allows wildcards to be used in the policy. For example, to allow any website on any port hosted in the "eur.ad.sag" domain, you can specify:

```
<init-param> <param-name>policy</param-name> <param-value>frame-ancestors
http://*.eur.ad.sag:* 'self' </param-value> </init-param>
```

8.2.12 Handle personal data in log files

For some actions MashZone NextGen tracks the user ID, IP address, email ID and full name of the executor. This data is used to analyze and fix potential problems that occur during the system operation. The data is also stored after deletion of a user account. This ensures that no past user data is lost and all user data are accounted for in future audits.

You can remove this personal data from your MashZone NextGen system for General Data Protection Regulation (page 445) (GDPR) compliance.

To remove the relevant personal user data, you must delete the corresponding log files from the MashZone NextGen system.

The relevant log files are stored in the following directories.

- IP address and user ID

```
<MashZone NextGen Installation>\apache-tomcat\logs\...
```

For example, localhost_access_log.2018-05-24

- Username

```
MashZone NextGen Installation>\apache-tomcat\logs\MashZone-AuditLog.log
```

- IP Address

```
MashZone NextGen Installation>\apache-tomcat\logs\wrapper.log
```

Warning

If you delete the log files, all logged data is lost and cannot be restored.

Procedure

In the Windows® Explorer, go to the directories mentioned above and delete all relevant log files. The user data is deleted from the MashZone NextGen system.

8.3 Getting Started with the MashZone NextGen Server

You install MashZone NextGen using the Software AG Installer. See the Installing Software AG Products guide for instructions.

The post-installation tasks you must complete to allow users to start working with MashZone NextGen include.

Procedure

1. Start the MashZone NextGen. See Start and Stop the MashZone NextGen Server (page 113) for instructions.
2. Login to MashZone NextGen:
 - a. Open MashZone NextGen in a browser at `http://localhost:8080/mashzone`.
If you used a different port number when you installed MashZone NextGen or the MashZone NextGen Server is running on a different host, change the domain and port number appropriately.
 - b. Use the credentials for the default administrator account:
User name = **Administrator**
Password = **manage**
3. If you are using the default MashZone NextGen User Repository rather than an LDAP Directory to manage users and groups for MashZone NextGen, it is a good practice to change the password for this default administrator account.
 - a. Open your profile from the MashZone NextGen Hub menu bar and click My Password.
 - b. Enter your new password and confirm this.
 - c. Then click Change Password.

If you will use your LDAP Directory as the MashZone NextGen User Repository, this default account is disabled once LDAP configuration is complete.
4. Set up a robust database to use as the MashZone NextGen Repository.
MashZone NextGen is installed with Derby as an embedded database hosting the MashZone NextGen Repository for trial purposes only. The default Derby database should **not** be used for serious development environments or for staging or production.
See Move the MashZone NextGen repository to a robust database solution (page 116) for instructions.

5. If you want MashZone NextGen to use your LDAP Directory as the repository for user and group information, you must update configuration. See [Integrate Your LDAP Directory with MashZone NextGen](#) (page 126) for instructions.
6. Configure the Event Service. See [Event Service Configuration and Administration](#) (Seite 221) for instructions.
7. If you have also installed Terracotta BigMemory and received your BigMemory license, add this license to MashZone NextGen and configure MashZone NextGen to work with BigMemory. See [Manage Licenses for MashZone NextGen and BigMemory and Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores](#) (page 170) for instructions.

8.3.1 Additional MashZone NextGen System and Software Requirements

- For basic requirements to install MashZone NextGen, see the [System Requirements for Software AG Products](#) guide.

8.3.1.1 Additional Recommendations for MashZone NextGen

In addition to the basic recommendations in the [System Requirements for Software AG Products](#) guide, you should also consider the following recommendations for MashZone NextGen:

- A robust, compatible database to host the MashZone NextGen Repository is required.
Important: The MashZone NextGen repository is initially installed in a Derby database suitable only for trial purposes. For proof-of-concept, development or production uses, move the repositories to a robust and compatible solution. See [System Requirements for Software AG Products](#) for more information.
- Architecture and memory requirements or recommendations include:
 - 64-bit architecture
 - 2G minimum of memory if only small to medium datasets are involved
 - 4G minimum of memory if large datasets are involved**Important:** Actual memory and disk space requirements vary widely based on load, throughput, performance and other requirements unique to your environment. Please contact your Software AG representative for more information.
- To ensure a secure connection between MashZone NextGen server and client it is recommended to operate your MashZone NextGen system via HTTPS as communication protocol. You can configure your application server accordingly after you have installed MashZone NextGen. See [Configure HTTPS and Certificate Stores in the Application Server](#) (page 162) for details.

8.3.2 What is Installed with MashZone NextGen

MashZone NextGen initially installs these WAR files:

WAR	Server and/or Application
presto.war	MashZone NextGen Server

MashZone NextGen also installs the following additional software:

- Apache's Tomcat Servlet Container, version 8.5.15
- Derby Database, version 10.5.3.0.

Important: The MashZone NextGen repository is initially installed in a Derby database suitable only for trial purposes. For proof-of-concept, development or production uses, move the repositories to a robust and compatible solution. See Move the MashZone NextGen repository to a robust database solution (page 116) for details.

8.3.3 Start and Stop the MashZone NextGen Server

Most MashZone NextGen widgets depend on the MashZone NextGen Server.

The MashZone NextGen Server depends on the MashZone NextGen Repository.

Startup and shutdown of MashZone NextGen Server does not automatically start and stop the integrated event service. Instead, the event service must be started separately, if required.

8.3.3.1 Start the MashZone NextGen Event Service

You can manually start the integrated Event Service, if required.

There are two ways to start the integrated Event Service for Windows systems.

On Windows Server operating systems MashZone NextGen Event Service must be started as Administrator.

Procedure

1. Run Start MashZone NextGen Event Service <version> in the program group Software AG > Start Servers of the Windows start menu.

To run MashZone NextGen Event Service as administrator, right click **Start MashZone NextGen Event Service <version>** and select **Run as administrator**.

2. Enter the following command in a command window.

c:>MashZoneNG-install/rtbs/bin/startup.bat

Starting startup.bat from the file system using Windows Explorer does not work.

For Linux, Mac OS X or UNIX systems, open a new terminal window and move to the **% cd MashZoneNG-install/rtbs/bin** folder and enter the **% startup.sh** command.

8.3.3.2 Start the MashZone NextGen Server

Procedure

1. If the MashZone NextGen Repository has been moved from the default Derby database and they are not already running, manually start these databases following the instructions for their host database.
2. Do one of the following to start the MashZone NextGen Server:
 - a. For Windows systems, either:

From the Start menu, select Software AG > Start Servers > Start MashZone NextGen version.

Enter this command in a command window:

```
c:>MashZoneNG-install/apache-tomcat/bin/startup.bat
```

On Windows Server operating systems MashZone NextGen Server must be started as Administrator. To run MashZone NextGen Server as Administrator, right click on the **Start MashZone NextGen version** shortcut and select **Run as administrator**.

Starting **startup.bat** from the file system using Windows Explorer does not work.
 - b. For Linux, Mac OS X or UNIX systems, open a new terminal window and move to this folder:

```
% cd MashZoneNG-install/apache-tomcat/bin
```

Then enter this command:

```
% startup.sh
```
3. Open the MashZone NextGen at **http://app-server:port/mashzone** and log in.

You can now access all the MashZone NextGen tools: Feed Editor, Dashboard Editor, Dashboard Viewer and the Admin Console.

8.3.3.3 Stop the MashZone NextGen Event Service

You can manually stop the integrated Event Service, if required.

There are two ways to stop the integrated Event Service for Windows systems.

Procedure

1. Run Stop MashZone NextGen Event Service <version> in the program group Software AG > Stop Servers of the Windows start menu.
2. Enter the following command in a command window.

```
c:>MashZoneNG-install/rtbs/bin/shutdown.bat
```

For Linux, Mac OS X or UNIX systems, open a new terminal window and move to the **% cd MashZoneNG-install/rtbs/bin** folder and enter the **% shutdown.sh** command.

8.3.3.4 Stop the MashZone NextGen Server

Procedure

1. Do one of the following:

For Windows systems, either:

From the Start menu, select **Software AG > Stop Servers > Stop MashZone NextGen**.

Enter this command in a command window:

```
c:>MashZoneNG-install/apache-tomcat/bin/shutdown.bat
```

For Linux, Mac OS X or UNIX systems, open a new terminal window and move to this folder:

```
% cd MashZoneNG-install/apache-tomcat/bin
```

Then enter this command:

```
% shutdown.sh
```

2. If the MashZone NextGen Repository has been moved from the default Derby database, you can also choose to stop this database. See documentation for the host database for instructions.

8.3.4 Startup Considerations

When the MashZone NextGen Repository is hosted in a robust database solution, it must be started before the MashZone NextGen Server for a successful startup. With the default Derby database, the MashZone NextGen Repository runs as an **embedded** database that is automatically started with the MashZone NextGen Server.

In environments where your application server is started automatically with the host, this can create timing errors. You may need to stop and restart the MashZone NextGen MashZone NextGen Server after the MashZone NextGen Repository has been restarted.

If you host the MashZone NextGen Repository in a MySQL or Oracle database, you may also be able to have the database start automatically with the host.

8.3.5 Manage Licenses for MashZone NextGen and BigMemory

To use MashZone NextGen a license is required.

If you are using BigMemory features that require this, you also need to make your BigMemory license available to the MashZone NextGen Server and/or the Integrated MashZone Server. See BigMemory for Caching, Connections and MashZone NextGen Analytics (page 167) for features that require this additional license.

You can apply licenses when you install MashZone NextGen, or you can install and use MashZone NextGen without a license for a trial period of 30 days. If you purchase MashZone NextGen after

installation, you must manually apply the MashZone NextGen license. If needed, you can also manually apply a BigMemory license.

When MashZone NextGen runs with a READ ONLY license, all tools to create and edit data feeds and dashboards are unavailable.

Procedure

1. Save the attached license file(s) from the email(s).
2. For MashZone NextGen licenses, copy the MashZoneNGLicense.xml file into the MashZoneNG-install/apache-tomcat/conf folder.
If MashZone NextGen is deployed in a cluster, copy the license file to this folder in every cluster member.
3. If a BigMemory license is required:
 - a. Copy the license file terracotta.key to the MashZoneNG-install/apache-tomcat/conf folder.
If MashZone NextGen is deployed in a cluster, you must copy this file to every cluster member.
 - b. Add the following Java system property to the MashZone NextGen server configuration file <MashZone NextGen installation>/apache-tomcat/conf/wrapper.conf:
wrapper.java.additional.<n+1>=-Dcom.tc.productkey.path=MashZoneNG-install/apache-tomcat/conf/terracotta.key
Where n is the number of last additional Java parameter.
If MashZone NextGen is deployed in a cluster, you must update the server configuration files for every cluster member.
4. Restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.3.6 Move the MashZone NextGen repository to a robust database solution

The MashZone NextGen repository is initially installed in a Derby database suitable only for trial purposes. For proof-of-concept, development or production uses, move the repositories to a robust and compatible solution.

You can host the MashZone NextGen repository in any database that MashZone NextGen supports. See Additional MashZone NextGen System and Software Requirements (page 112) in System Requirements for details.

You can move the repository to one of the following databases:

- **Microsoft SQL Server:** see Move MashZone NextGen repository to Microsoft SQL Server (page 118) for instructions.
- **MySQL:** see Move the MashZone NextGen repository to MySQL (page 120) for instructions.
- **Oracle:** see Move the MashZone NextGen repository to Oracle (page 122) for instructions.

- **PostGres:** see Move the MashZone NextGen repository to PostGres (page 124) for instructions.

8.3.6.1 Troubleshooting Connections to the MashZone NextGen Repository

The most common problem when the MashZone NextGen Server server does not restart successfully after you move the MashZone NextGen Repository to a new database is that it cannot connect to the MashZone NextGen Repository. To verify that this is the problem:

- Open the MashZone NextGen Server log file **prestoserver.log** in your web application server's log directory. For Tomcat, this is:
web-apps-home/logs/prestoserver.log
- Check for a log entry for **Cannot create PoolableConnectionFactory** near the end of the file. This error indicates the MashZone NextGen Server could not successfully connect to the new database.

Common causes for this error include:

- The database hosting the MashZone NextGen Repository is not running.
- If this is true, start the MashZone NextGen Repository and verify that it is up. Then restart the MashZone NextGen Server and confirm that this starts successfully.
- There are one or more firewalls between the MashZone NextGen Repository and the MashZone NextGen Server that are not configured to allow this connection.
This can only happen when the database for the MashZone NextGen Repository is hosted on a different server than the MashZone NextGen Server.
- Update the firewall configuration to allow this connection. Then restart the MashZone NextGen Server and confirm that this starts successfully.
- The URL or other connection configuration that you entered in MashZone NextGen for the MashZone NextGen Repository is incorrect.
- To correct an error in this case, edit the resource properties for the MashZone NextGen Repository in the MashZoneNG-install/apache-tomcat/conf/context.xml file.
- Then restart the MashZone NextGen Server and confirm that this starts successfully.
- Port or connection configuration for the database is not set up properly to allow connections from the MashZone NextGen Server. See documentation for your database for more information.

8.3.6.2 Move the MashZone NextGen repository to Microsoft SQL Server

Procedure

1. If you are using your LDAP Directory as the MashZone NextGen User Repository, make sure that at least one user in your LDAP Directory has administrator privileges for MashZone NextGen before you move the MashZone NextGen Repository. See Grant User Access to MashZone NextGen with Built-in Groups (page 104) for instructions.

When the MashZone NextGen User Repository is your LDAP Directory, the default administrator account (**Administrator** user) is disabled.

2. If you are hosting the MashZone NextGen Repository or MashZone Repository in a new database, create the database following SQL Server documentation. Keep the following points in mind:

Make sure this database is supported by MashZone NextGen. See Additional MashZone NextGen System and Software Requirements (page 112) for details.

The JTDS driver and the original Microsoft driver are available for Microsoft SQL Server. You must make different settings depending on the driver type selected. For details see the following steps.

If you want MashZone NextGen to support international characters in meta-data for artifacts, make sure the database uses the UTF-16 character encoding and case insensitive collation. See documentation for your database for specific instructions.

It is a best practice to require passwords for every database account that can access the MashZone NextGen Repository.

If you do not use the default **dbo** schema, you have to specify the name of the used schema (value="schema_name") in the bean PMF in the **rdsApplicationContext.xml** file.

The file is located in <MashZone NextGen

installation>\apache-tomcat\webapps\mashzone\WEB-INF\classes\.

```
<bean id="pmf" class="com.jackbe.jbp.sas.rds.impl.jdo
.PersistenceManagerFactoryBean"> <property name="lifecycleListener"
ref="entityLifecycleLsnr"/> <property name="dataSource" ref="dataSource" />
<property name="configLocation" value="classpath:datanucleus.properties"/>
<!-- overwrite settings in configLocation file --> <property
name="jdoProperties"> <map> <entry key="datanucleus.ConnectionFactoryName"
value="java:comp/env/MashzoneNextGenRepository"/> <entry
key="datanucleus.storeManagerType" value="rdbms"/> <entry
key="datanucleus.mapping.Schema" value="schema_name"/> </map> </property>
</bean>
```

3. Start the database that will become host to the MashZone NextGen Repository, if it is not already up.
4. Using the SQL tool for the database that will be host, add MashZone NextGen Repository tables with the scripts shown below from the corresponding folder in MashZoneNG-install/prestorepository/mssqldb:

createDBTables.txt for MetaData and the default User Repository

createSnapsTables.sql for Snapshots

createSchedulerTables.sql for Scheduler

This folder also contains scripts to drop the corresponding MashZone NextGen Repository tables, if needed.

5. Copy the JAR file for the JDBC driver for your database to the following folder for each MashZone NextGen Server that uses this MashZone NextGen Repository:
MashZoneNG-install/apache-tomcat/lib
6. Replace the JAR for the MashZone NextGen Repository:
 - a. Remove the
web-apps-home/mashzone/WEB-INF/lib/jackbe-presto-rds-postgresql-derby.jar JAR file for each MashZone NextGen Server that uses this MashZone NextGen Repository. You can delete this JAR or simply move it to a folder that is not in the classpath for the application server that hosts MashZone NextGen.
 - b. Copy this JAR file:
MashZoneNG-install/prestorepository/jackbe-presto-rds-oracle-mysql-mssql.jar
To the web-apps-home/mashzone/WEB-INF/lib folder.
7. Update snapshot scheduler configuration for the MashZone NextGen Server:
 - a. In the text editor of your choice, open the applicationContext-scheduler.xml file in the webapps-home/mashzone/WEB-INF/classes/ folder for the MashZone NextGen Server.
 - b. Find the bean for org.springframework.scheduling.quartz.SchedulerFactoryBean in default profile.
 - c. Update the org.quartz.jobStore.driverDelegateClass property to the org.quartz.impl.jdbcjobstore.MSSQLDelegate appropriate delegate for this database:
 - d. Save this change.
 - e. If this is a clustered environment, copy the updated applicationContext-scheduler.xml configuration file to each MashZone NextGen Server in the cluster.
8. Open the MashZoneNG-install/apache-tomcat/conf/context.xml configuration file in the text editor of your choice.
9. For the MashZone NextGen Repository, edit the <Resource> element with an ID of MashzoneNextGenRepository and:
 - a. Update the JDBC driver, URL and credential properties:

Example for jTDS driver

```
<Resource name="MashzoneNextGenRepository" auth="Container"
type="javax.sql.DataSource" maxTotal="200" maxIdle="30"
maxWaitMillis="10000" username="app" password="app"
driverClassName="net.sourceforge.jtds.jdbc.Driver"
url="jdbc:jtds:sqlserver://host-name:port/database" />
```

The JTA managed property **must** be **false**.

Example for original Microsoft driver

```
<Resource name="MashzoneNextGenRepository" auth="Container"
type="javax.sql.DataSource" maxTotal="200" maxIdle="30"
maxWaitMillis="10000" username="app" password="app"
```

```
driverClassName="com.microsoft.sqlserver.jdbc.SQLServerDriver"  
url="jdbc:sqlserver://host-name:port;databaseName=your_database" />
```

- b. If needed, update optional properties. See Tomcat Datasource Properties (<http://commons.apache.org/proper/commons-dbcp/configuration.html>) for a complete list of optional properties and information on defaults.

Some common properties you may need to set include:

validationQuery = select 1

Common tuning properties for connections pools. See Tuning the MashZone NextGen Repository Connection Pool (page 267).

10. Save your changes to this file.

If the MashZone NextGen Server does not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.

11. Restart the MashZone NextGen Server to apply these changes.

If the MashZone NextGen Server does not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.

8.3.6.3 Move the MashZone NextGen repository to MySQL

Procedure

1. If you are using your LDAP Directory as the MashZone NextGen User Repository, make sure that at least one user in your LDAP Directory has administrator privileges for MashZone NextGen before you move the MashZone NextGen Repository. See Grant User Access to MashZone NextGen with Built-in Groups (page 104) for instructions.

When the MashZone NextGen User Repository is your LDAP Directory, the default administrator account (**Administrator** user) is disabled.

2. If you are hosting the MashZone NextGen Repository in a new database, create the database following MySQL documentation (<http://dev.mysql.com/doc/>). Keep the following points in mind:

Make sure this database is supported by MashZone NextGen. See Additional MashZone NextGen System and Software Requirements (page 112) for details.

If you want MashZone NextGen to support international characters in meta-data for artifacts, set the character encoding and collation to UTF-8 when you create the database. See documentation for your database for specific instructions.

For medium or larger MySQL databases that will host the MashZone NextGen Repository, you should increase the maximum allowed packet size, which defaults to 1MB, for the database.

3. Start the database that will become host to the MashZone NextGen Repository, if it is not already up.

4. Using the SQL tool for the database that will be host, add MashZone NextGen Repository tables with the scripts shown below from the corresponding folder in MashZoneNG-install/prestorepository/mysql:
 - a. Create a Database to hold the MashZone NextGen Repository tables. See the file createDB.txt for an example.
 - b. Create a User with rights to the database created in step a. See the file createUser.txt for an example.
 - c. Connect to the MySQL database created in step a using a SQL tool (for example MySQL command line client) using the user created in step b. See the comments in the file createDBTables.txt for examples.
 - d. Execute the statements in the file createDBTables.txt to create the tables using the SQL tool (for example, use the MySQL source command: "source /path/to/createDBTables.txt").
 - e. Execute the statements in the file createSchedulerTables.sql (for example: "source /path/to/createSchedulerTables.sql").
 - f. Execute the statements in the file createSnapsTables.sql (for example: "source /path/to/createSnapsTables.sql").

This folder also contains scripts to drop the corresponding MashZone NextGen Repository tables, if needed.

5. Replace the JAR for the MashZone NextGen Repository:
 - a. Remove the MashZoneNG-install/apache-tomcat/mashzone/WEB-INF/lib/jackbe-presto-rds-postgresql-derby.jar JAR file for each MashZone NextGen Server that uses this MashZone NextGen Repository. You can delete this JAR or simply move it to a folder that is not in the classpath for the application server that hosts MashZone NextGenMashZone NextGen.
 - b. Copy this JAR file:
MashZoneNG-install/prestorepository/jackbe-presto-rds-oracle-mysql-mssql.jar
To the web-apps-home/mashzone/WEB-INF/lib folder.
6. Copy the MySQL JDBC driver jar file to MashZoneNG-install/apache-tomcat/lib.
7. Open the MashZoneNG-install/apache-tomcat/conf/context.xml configuration file in the text editor of your choice.
8. For the MashZone NextGen Repository, edit the <Resource> element with an ID of MashzoneNextGenRepository and:
 - a. Update the JDBC driver, URL and credential properties:

```
name="MashzoneNextGenRepository" auth="Container"
type="javax.sql.DataSource" maxTotal="200" maxIdle="30"
maxWaitMillis="10000" username="app" password="app"
driverClassName="com.mysql.jdbc.Driver"
url="jdbc:mysql://host-name/databasename" />
```

For MySQL databases, it is **recommended** that you include the database name in data source URLs. If this information is omitted, testing the data source fails and may also cause errors with access to stored procedures.

The JTA managed property **must** be **false**.

- b. If needed, update optional properties. See Tomcat Datasource Properties (<http://commons.apache.org/proper/commons-logging/configuration.html>) for a complete list of optional properties and information on defaults.

Some common properties you may need to set include:

validationQuery = select 1 from dual

Common tuning properties for connections pools. See Tuning the MashZone NextGen Repository Connection Pool (page 267).

9. Start the MashZone NextGen Server to apply these changes. This also starts the MashZone Server.

If the MashZone NextGen Server does not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.

8.3.6.4 Move the MashZone NextGen repository to Oracle

Procedure

1. If you are using your LDAP Directory as the MashZone NextGen User Repository, make sure that at least one user in your LDAP Directory has administrator privileges for MashZone NextGen before you move the MashZone NextGen Repository. See Grant User Access to MashZone NextGen with Built-in Groups (page 104) for instructions.

When the MashZone NextGen User Repository is your LDAP Directory, the default administrator account (Administrator user) is disabled.

2. If you are hosting the MashZone NextGen Repository in a new database, create the database following Oracle documentation (<http://www.oracle.com/technetwork/database/enterprise-edition/documentation/index.html>).

Make sure this database is supported by MashZone NextGen. See Additional MashZone NextGen System and Software Requirements (page 112) for details.

If you want MashZone NextGen to support international characters in meta-data for artifacts, set the character encoding to AL32UTF8 when you create the database. See documentation for your database for specific instructions.

It is a best practice to require passwords for every database account that can access the MashZone NextGen Repository.

3. Start the database that will become host to the MashZone NextGen Repository, if it is not already up.

4. Using the SQL tool for the database that will be host, add MashZone NextGen Repository tables with the scripts shown below from the corresponding folder in MashZoneNG-install/prestorepository/oracledb:

createDB.txt

createDBTables.txt for MetaData and the default User Repository

createSnapsTables.sql for Snapshots

createSchedulerTables.sql for Scheduler

This folder contains other scripts to drop the corresponding MashZone NextGen Repository tables.

5. Replace the JAR for the MashZone NextGen Repository:
 - a. Remove the web-apps-home/mashzone/WEB-INF/lib/jackbe-presto-rds-postgresql-derby.jar JAR file for each MashZone NextGen Server that uses this MashZone NextGen MashZone NextGen Repository. You can delete this JAR or simply move it to a folder that is not in the classpath for the application server that hosts MashZone NextGen.
 - b. Copy this JAR file:
MashZoneNG-install/prestorepository/jackbe-presto-rds-oracle-mysql-mssql.jar
To the web-apps-home/mashzone/WEB-INF/lib folder.
6. Copy the JAR file for the JDBC driver for your database to the following folder for each MashZone NextGen Server that uses this MashZone NextGen Repository:
MashZoneNG-install/apache-tomcat/lib
7. Update snapshot scheduler configuration for the MashZone NextGen Server:
 - a. In the text editor of your choice, open the applicationContext-scheduler.xml file in the webapps-home/mashzone/WEB-INF/classes/ folder for the MashZone NextGen Server.
 - b. Find the bean for org.springframework.scheduling.quartz.SchedulerFactoryBean in default profile.
 - c. Update the org.quartz.jobStore.driverDelegateClass property to the org.quartz.impl.jdbcjobstore.oracle.OracleDelegate delegate.

The configuration would now look like:

```
... <bean id="scheduler"> <property
name="applicationContextSchedulerContextKey">
<value>applicationContext</value> </property> <property
name="quartzProperties"> <props> <prop
key="org.quartz.scheduler.instanceId">AUTO</prop> <prop
key="org.quartz.jobStore.class"> org.quartz.impl.jdbcjobstore.
JobStoreTX</prop> <prop key="org.quartz.jobStore.tablePrefix">QRTZ_</prop>
<prop key="org.quartz.jobStore.driverDelegateClass"> org.quartz.impl.
jdbcjobstore.oracle.OracleDelegate</prop> <prop
key="org.quartz.jobStore.dataSource">schedulerDS</prop> ... </props>
</property> </bean> ...
```

- d. Save this change.
- e. If this is a clustered environment, copy the updated applicationContext-scheduler.xml configuration file to each MashZone NextGen Server in the cluster.

8. Open the MashZoneNG-install/apache-tomcat/conf/context.xml configuration file in the text editor of your choice.
9. For the MashZone NextGen Repository, edit the <Resource> element with an ID of MashzoneNextGenRepository and:
 - a. Update the JDBC driver, URL and credential properties:

```
<Resource name="MashzoneNextGenRepository" auth="Container"
type="javax.sql.DataSource" maxTotal="200" maxIdle="30"
maxWaitMillis="10000" username="app" password="app"
driverClassName="oracle.jdbc.OracleDriver"
url="jdbc:oracle:driverType:@host-name:port:dbname" />
```

The JTA managed property **must** be **false**.
 - b. If needed, update optional properties. See Tomcat Datasource Properties (<http://commons.apache.org/proper/commons-dbc/p/configuration.html>) for a complete list of optional properties and information on defaults.
Some common properties you may need to set include:
validationQuery = select 1 from dual
Common tuning properties for connections pools. See Tuning the MashZone NextGen Repository Connection Pool (page 267).
10. Save your changes to this file.
If the MashZone NextGen Server does not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.
11. Restart the MashZone NextGen Server to apply these changes.
If the MashZone NextGen Server does not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.

8.3.6.5 Move the MashZone NextGen repository to PostGres

Procedure

1. If you are using your LDAP Directory as the MashZone NextGen User Repository, make sure that at least one user in your LDAP Directory has administrator privileges for MashZone NextGen before you move the MashZone NextGen Repository. See Grant User Access to MashZone NextGen with Built-in Groups (page 104) for instructions.
When the MashZone NextGen User Repository is your LDAP Directory, the default administrator account (**Administrator** user) is disabled.
2. If you are hosting the MashZone NextGen Repository or MashZone Repository in a new database, create the database following PostgreSQL documentation (<http://www.postgresql.org/docs/>). Keep the following points in mind:
Make sure this database is supported by MashZone NextGen. See Additional MashZone NextGen System and Software Requirements (page 112) for details.

If you want MashZone NextGen to support international characters in meta-data for artifacts, set the character encoding to UTF8 when you create the database. See documentation for your database for specific instructions.

It is a best practice to require passwords for every database account that can access the MashZone NextGen Repository.

When you initialize the Postgres database that will host the MashZone NextGen Repository, you may need to specifically set the locale used by the database to ensure case-insensitive sorting.

3. Start the database that will become host to the MashZone NextGen Repository, if it is not already up.
4. Using the SQL tool for the database that will be host, add MashZone NextGen Repository tables with the scripts shown below from the corresponding folder in MashZoneNG-install/prestorepository/postgresdb:

createDBTables.txt for MetaData and the default User Repository

createSnapsTables.sql for Snapshots

createSchedulerTables.sql for Scheduler

There are also scripts to drop the corresponding MashZone NextGen Repository tables in these folders, if needed.

5. Update snapshot scheduler configuration for the MashZone NextGen Server:
 - a. In the text editor of your choice, open the applicationContext-scheduler.xml file in the webapps-home/mashzone/WEB-INF/classes/ folder for the MashZone NextGen Server.
 - b. Find the bean for org.springframework.scheduling.quartz.SchedulerFactoryBean in default profile.
 - c. Update the org.quartz.jobStore.driverDelegateClass property to the org.quartz.impl.jdbcjobstore.PostgreSQLDelegate delegate.

The configuration would now look like:

```
... <bean id="scheduler" > <property
name="applicationContextSchedulerContextKey">
<value>applicationContext</value> </property> <property
name="quartzProperties"> <props> <prop
key="org.quartz.scheduler.instanceId">AUTO</prop> <prop
key="org.quartz.jobStore.class"> org.quartz.impl
.jdbcjobstore.JobStoreTX</prop> <prop
key="org.quartz.jobStore.tablePrefix">QRTZ_</prop> <prop
key="org.quartz.jobStore.driverDelegateClass"> org.quartz
.impl.jdbcjobstore.PostgreSQLDelegate</prop> <prop
key="org.quartz.jobStore.dataSource">schedulerDS</prop> ... </props>
</property> </bean> ...
```

- d. Save this change.
 - e. If this is a clustered environment, copy the updated applicationContext-scheduler.xml configuration file to each MashZone NextGen Server in the cluster.
6. Copy the JAR file for the JDBC driver for your database to the following folder for each MashZone NextGen Server that uses this MashZone NextGen Repository:
MashZoneNG-install/apache-tomcat/lib.

7. Open `rdsApplicationContext.xml` under `MashZoneNG-install/apache-tomcat/classes` and add the following keys to:

```
<property name="jdoProperties"> ... <map> ... <entry key="javax.jdo.mapping.Schema" value="public"/> <entry key="datanucleus.identifier.case" value="LowerCase"/> ... </map> </property>
```
8. If you are using PostgreSQL version 9.x please open `postgresql.conf` under `PostgreSQL-install/9.x/data` and un-comment the following property and make sure it is set to off: `standard_conforming_strings = off`.
9. Open the `MashZoneNG-install/apache-tomcat/conf/context.xml` configuration file in the text editor of your choice.
10. For the MashZone NextGen Repository, edit the `<Resource>` element with an ID of `MashzoneNextGenRepository` and:
 - a. Update the JDBC driver, URL and credential properties:

```
name="MashzoneNextGenRepository" auth="Container" type="javax.sql.DataSource" maxTotal="200" maxIdle="30" maxWaitMillis="10000" username="app" password="app" driverClassName="org.postgresql.Driver" url="jdbc:postgresql://host-name:port/databasename" />
```
 - b. The JTA managed property **must** be **false**.

If needed, update optional properties. See Tomcat Datasource Properties (<http://commons.apache.org/proper/commons-dbcp/configuration.html>) for a complete list of optional properties and information on defaults.

Some common properties you may need to set include:

validationQuery = select 1

Common tuning properties for connections pools. See Tuning the MashZone NextGen Repository Connection Pool (page 267).
11. Save your changes to this file.

If the MashZone NextGen Server does not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.
12. Restart the MashZone NextGen Server to apply these changes. This also restarts the MashZone Server.

If the MashZone NextGen Server wedoes not start up successfully, see Troubleshooting Connections to the MashZone NextGen Repository (page 117) for suggestions.

8.3.7 Integrate Your LDAP Directory with MashZone NextGen

In many cases, users and authentication information for an organization is defined in an existing LDAP Directory. You can configure MashZone NextGen to use your LDAP Directory as the source for user and group information.

See the System Requirements for Software AG Products guide for information on MashZone NextGen support for specific LDAP Directory solutions.

Procedure

1. If the MashZone NextGen Server is not yet started, start MashZone NextGen. See Start and Stop the MashZone NextGen Server (page 113) for instructions.
2. Change MashZone NextGen configuration to use LDAP as the authentication provider.
 - a. Edit the `userRepositoryApplicationContext.xml` file in the `MashZoneNG-config` folder with any text editor.

This folder may be in the default location or in an external location. See Setting Up an External MashZone NextGen Configuration Folder (page 213) for more information.

- b. Remove the comment markers around this statement: `<import resource="/userRepositoryApplicationContext-ldap.xml">`.
- c. Comment out this statement: `<import resource="/userRepositoryApplicationContext-jdbc.xml">` property.

You cannot use both default authentication and LDAP authentication.

The configuration should look something like this:

```
<beans> <!-- Choose between the internal JDBC repository and LDAP comment/uncomment
these two import statements --> <import
resource="/userRepositoryApplicationContext-ldap.xml"> <!-- <import
resource="/userRepositoryApplicationContext-jdbc.xml"> --> ... </beans>
```

3. Change MashZone NextGen configuration for the user attribute provider.
 - a. If it is not already open, edit the `userRepositoryApplicationContext.xml` file in the `MashZoneNG-config` folder with any text editor.

This folder may be in the default location or in an external location. See Setting Up an External MashZone NextGen Configuration Folder (page 213) for more information.
 - b. Find the `userAttributeProvider` bean:

```
<bean id="userAttributeProvider" > ...
```
 - c. Remove comment markers around the `LdapAttributeProvider` bean reference in the providers property list.

The configuration should now look something like this:

```
<bean id="userAttributeProvider" > <property name="providers"> <list> <ref
bean="ldapAttributeProvider"/> <ref bean="internalUserAttributeProvider"/>
</list> </property> </bean>
```

- d. Save your changes to this file.

Do **not** restart the MashZone NextGen Server until all other LDAP configuration steps have been completed.

4. Define configuration in the Admin Console in MashZone NextGen Hub for:
 - Connections to the LDAP Directory. See Defining LDAP Connection Configuration (page 128).
 - Authentication mechanisms. See Defining the Authentication Scheme (page 128).
 - Authorization mechanisms. See Defining the Authorization Scheme (page 129).

All user and group queries used in MashZone NextGen applications. See [Enabling MashZone NextGen Application Queries for All LDAP Users or Groups for Permissions](#) (page 131).

5. Add the built-in MashZone NextGen user groups to LDAP as new groups and assign at least one user as a MashZone NextGen administrator.

See [Grant User Access to MashZone NextGen with Built-in Groups](#) (page 104) for instructions.

6. Restart the MashZone NextGen Server.

MashZone NextGen now uses LDAP as the user repository. You can now login using the user account assigned in earlier steps as a MashZone NextGen administrator.

To grant access to other users, add them to an appropriate built-in MashZone NextGen user group in LDAP. See [Grant User Access to MashZone NextGen with Built-in Groups](#) (page 104) for instructions.

8.3.7.1 Defining LDAP Connection Configuration

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.
3. Set these properties for your LDAP Directory:
 - LDAP URL** = the URL to your LDAP directory. For example:
ldap://localhost:389/dc=somecompany,dc=com
 - Directory User Name** = the distinguished name for the user account to connect to this LDAP Directory. This **must** be a privileged account. For example:
uid=admin,ou=system
 - Directory Password** = the password for the user to connect to this LDAP Directory.
4. Change any advanced options (see [Defining the Authentication Scheme](#) (page 128), [Defining the Authorization Scheme](#) (page 129) and [Enabling MashZone NextGen Application Queries for All LDAP Users or Groups for Permissions](#) (page 131)).
5. Save your changes.

8.3.7.2 Defining the Authentication Scheme

Authentication against LDAP determines if a distinguished name exists for a user. This searches for a user entry based on a specific username. Search-based authentication works, for example, if user names are users' email addresses.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.

3. Click **Advanced Options**.
4. Set these properties in the Authentication Properties section:

User Search Base = the base context for a user search in authentication. This produces a list of all users which is filtered with a combination of the User Search Filter and User Search Subtree properties to authenticate a user. For example:

ou=users,ou=system

User Search Filter = the relative filter to apply to search for users during authentication. The variable **{0}** is replaced with the user's username from login.

This filter is based from the context defined in User Search Base. For example:

email={0}

This attribute **must** be the same attribute used in the **User ID Attribute Name** property.

User Search Subtree = set this option if the search should be recursive through all levels of the Directory under the search base. If you clear this option, search only checks direct children of the search base.

Use LDAP VLV Control for Sorting and Paging = this option is set by default to allow MashZone NextGen to use **virtual list views** (VLV) to paginate and sort LDAP search results. Most LDAP directories support VLV, so in most cases you can leave this option set. If your LDAP directory logs errors for "unsupported search control", you can use this option to turn VLV off.

User ID Attribute Name = the LDAP attribute that contains the username that users login with. For example:

email

This value becomes the user ID for all further security contexts, unless the User ID Pattern property is also set.

User ID Pattern = a regular expression that is applied to user login names to extract the user ID for all further security contexts. This is only applied after authentication occurs.

8.3.7.3 Defining the Authorization Scheme

MashZone NextGen permissions are assigned to user groups or to individual users. To set up authorization when LDAP is the user repository, you must relate MashZone NextGen user groups to user groups in LDAP and define how users are assigned to groups in LDAP. User membership in LDAP groups can be defined by adding users to group entries or by adding group names to user entries, but **not** both.

In previous releases, MashZone NextGen user groups were called roles that could be implemented as user roles in LDAP instead of user groups. To use roles in LDAP for authorization in MashZone NextGen, please contact your Software AG representative for more information.

You **must** add the built-in MashZone NextGen groups that define basic permissions as groups in LDAP. You assign users to these built-in groups to assign basic MashZone NextGen permissions. Your existing LDAP groups can then be used in MashZone NextGen to

define permissions for specific dashboards, data feeds and aliases. For more information on authorization, see Authorization Policies and Permissions (page 104).

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.
3. Click **Advanced Options**.
4. If user membership is defined in group entries in your LDAP directory, set these properties:

Set the **Search Groups for User Membership** option.

Enter the beginning context for user group searches in the **Group Search Base** property.

This is combined with the User Group Search Filter to find LDAP groups to determine user membership in groups that may have MashZone NextGen permissions. For example:

ou=groups

Enter the filter to apply in group searches in the **User Group Search Filter** property.

This is combined with Group Search Base to find LDAP groups to determine user membership in groups that may have MashZone NextGen permissions. The variable **{0}** is replaced with the user's username from login. For example:

uniquemember={0}

Enter the LDAP attribute in group entries that identifies a group in the **Group Name Attribute** property.

This attribute contains the name of user groups that is used in MashZone NextGen permissions. The default value is the group common name:

cn

If you change this property, you **must** also update the **Group Name Pattern** property.

If group IDs in your LDAP Directory are not simple common names (see Group Name Attribute), enter a regular expression in **Group Name Pattern** to identify the built-in MashZone NextGen groups.

For example:

cn=(PRESTO_.*?)

MashZone NextGen expects specific names for the built-in groups that you add to your LDAP Directory. These values are defined in the common name of the group. This property allows MashZone NextGen to find the expected values for built-in groups, but use the full correct group names for the groups for your organization.

5. If user membership is defined solely in user entries, set these properties:

Clear the **Search Groups for User Membership** option.

Enter the name of the LDAP attribute in user entries that identifies the groups that users belong to in the **User Membership Attribute** property.

If group IDs in your LDAP Directory are not simple common names, enter a regular expression in **Group Name Pattern** to identify the built-in MashZone NextGen groups.

For example:

cn=(PRESTO_.*?)

MashZone NextGen expects specific names for the built-in groups that you add to your LDAP Directory. These values are defined in the common name of the group. This property allows MashZone NextGen to find the expected values for built-in groups, but use the full correct group names for the groups for your organization.

With these properties set, for example:

```
Search Groups for User Membership = true Group Search Base = ou= groups,ou=system
User Group Search Filter=uniquemember={0} Group Name Attribute = cn
```

And a username of **jwtalker**, MashZone NextGen would search all entries in **ou=groups** where **uniquemember=jwtalker**. The names for any of these groups would be the common name (cn) for the group entry.

If these properties were set instead:

```
Search Groups for User Membership = false User Membership Attribute = memberOf
```

The list of groups would consist of all values in the **memberOf** attribute in the **jwtalker** user entry.

This list of group names would be compared to the built-in MashZone NextGen groups and to groups with permissions for artifacts to determine the full set of permissions for **jwtalker**.

8.3.7.4 Enabling MashZone NextGen Application Queries for All LDAP Users or Groups for Permissions

MashZone NextGen queries the User Repository for user groups and users to enable you and other users to assign permissions for MashZone NextGen resources. To enable these queries you set properties in the Admin Console.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.
3. Click **Advanced Options**.
4. To enable queries for all users, set these properties:

User Search Base (in Authentication Properties) = the base context for a search for all users. This is used with the All Users Search Filter and Search Subtree For All Users properties to get a result. For example:

ou=People

This property is also used to search for users during authentication. Consider both uses before changing its value.

All Users Search Filter (in MashZone NextGen Queries) = the search filter, combined with User Search Base that is used to find all user entries. For example:

objectclass=inetOrgPerson

Ensure that the **objectclass=inetOrgPerson** attribute is set on the LDAP server.

5. **To support wildcard searches and** define the sort order for results, you must also define these properties:

Attributes Used in Wildcard Search (in MashZone NextGen Queries) = a list of LDAP attributes, separated by commas, to search in for wildcard searches. This defaults to:

cn,uid

User Sort By Attribute (in MashZone NextGen Queries) = the LDAP attribute that should be used to sort the results of wildcard searches. This defaults to:

cn

6. You must also define these properties so that Admin Console can display minimal user information:

User First Name Attribute (in MashZone NextGen Queries) = the LDAP attribute that holds users' first names.

User Last Name Attribute (in MashZone NextGen Queries) = the LDAP attribute that holds users' last names.

User Email Attribute (in MashZone NextGen Queries) = the LDAP attribute that holds users' email addresses.

7. To enable queries for LDAP groups that can be used to assign MashZone NextGen permissions:

Group Search Base (in Authorization Properties) = the beginning context, combined with Filter to Find All Groups for Roles to find all LDAP groups that can be used to assign MashZone NextGen permissions. For example:

ou=groups

This property is also used to search for MashZone NextGen permissions during authorization. Consider both uses before changing its value.

Filter to Find All Groups for Permissions = the search filter, combined with Group Search Base that is used to find all LDAP groups that may be used to assign MashZone NextGen permissions. For example:

objectclass=groupOfUniqueNames

Trouble shooting: If your LDAP user with role Presto_Administrator does not work, it might be helpful to stop MashZone NextGen first, deactivate and reactivate your LDAP connection in MashZone NextGen and then restart MashZone NextGen again.

8.3.8 Use the Default MashZone NextGen User Repository

MashZone NextGen authenticates users against a repository and retrieve authorization from the repository. This can be either a database or an LDAP Directory.

MashZone NextGen is installed with a default user repository within the MashZone NextGen Repository. You can use this initially as you explore MashZone NextGen or keep as the permanent source for user information, if desired. The default user repository also contains Default User Accounts (page 100) that you can use initially.

The MashZone NextGen repository is initially installed in a Derby database suitable only for trial purposes. For proof-of-concept, development or production uses, move the repositories to a robust and compatible solution.

No configuration is required to use the default user repository. If you use this default, you manage user and group information using functions in the Admin Console. See these topics for more information:

- Manage Users (page 133)
- Manage User Groups (page 134)
- Automatically Assign New Users to Groups (page 135)

8.3.8.1 Manage Users

If you are using the default MashZone NextGen User Repository, MashZone NextGen administrators can add users, assign them to groups to grant them permissions for various actions, and otherwise manage users in the Admin Console to. See Create Users (page 133) and Edit, Grant Permissions and other User Management Tasks (page 134) for instructions.

If you have configured MashZone NextGen to use your LDAP Directory as the User Repository, you manage users and groups and assign permissions in LDAP.

8.3.8.1.1 Create Users

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Users and Groups** tab and click **Users**.
3. Click **Add new user** and complete the following information:

User Name must be unique. This is the end-user's login name and primary key for the user record. User names cannot exceed 256 characters.

First Name and **Last Name** are optional.

Email

Password cannot exceed 50 characters. Confirm the password.

Valid characters for the User Name and Password fields are defined by the database you use for the default MashZone NextGen User Repository.

4. Set the **Active** option, if needed, and click **Add this User**.
5. Add another user or click **Cancel** to close this form.

The new user is now active in MashZone NextGen and has permissions as an authenticated user. They can also access any artifact that grants permissions to the All Authenticated Users group.

To give new users access to the MashZone NextGen to create artifacts or simply to work with artifacts from other users, you must add them to other groups, either manually or automatically.

For more information, see:

- Automatically Assign New Users to Groups (page 135)
- Edit, Grant Permissions and other User Management Tasks (page 134)

8.3.8.1.2 Edit, Grant Permissions and other User Management Tasks

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Users and Groups** tab and click **Users**.
3. Find the user in the list using Search or scrolling. Management options include.
4. Click **Reset password** and enter a new password.
5. Click  **Edit** to change the user's name, login name or email.
6. Click **Change user status** to activate the user or click **Change user status** to deactivate.
7. Click **Manage User Groups** to grant or manage permissions for this user.
 - a. Enter part of a group name, if needed, and click **Search** to see a list of user groups with permissions and the user groups/permissions currently assigned to this user.
 - b. Click a group from the left column to assign a user group/permission to this user. Click a group from the right pane to remove a user group/permission.
 - c. See Built-In MashZone NextGen User Groups and Permissions (page 105) for information on the permissions for the built-in user groups in MashZone NextGen.
 - d. Click **Save changes** to update groups and permissions for this user.
8. Click  **Delete** and confirm. Note that you cannot delete the default administrator account (the **admin** user) for MashZone NextGen.

8.3.8.2 Manage User Groups

Groups are used to grant permissions to users for specific actions in MashZone NextGen. If you are using the default MashZone NextGen User Repository, MashZone NextGen administrators add and manage user groups in the Admin Console. If you have configured MashZone NextGen to use your LDAP Directory as the User Repository, you define and manage user groups in your LDAP Directory.

User groups contain of a set of users. Groups also have one or more permissions assigned.

Built-in MashZone NextGen groups have a set of permissions predefined that allow users to work in MashZone NextGen. See Built-In MashZone NextGen User Groups and Permissions (page 105) for details.

Note: MashZone NextGen administrators and artifact owners automatically have permission to all artifacts or the artifacts they own respectively.

User-defined groups are groups that administrators add to MashZone NextGen. These groups are used to grant access permissions to group members for the specific artifacts that the group has been added to.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Users and Groups** tab and click **Groups**.
A list of groups displays. You can filter this list by entering part of a group name and clicking **Search**.
3. Click **Add new user group** and:
 - a. Enter a unique name in the **Group** field.
 - b. Click **Add this Group**.
 - c. Add more groups or click **Cancel** to close this form.
4. To delete a user group, click  **Delete** and confirm this.
Do **not** delete any of the built-in groups for MashZone NextGen: **Presto_Administrator**, **Presto_Developer**, **Presto_PowerUser**, **Presto_AuthenticatedUser** or **Presto_Guest**.

8.3.8.3 Automatically Assign New Users to Groups

If you are using the Default MashZone NextGen User Repository, you can automatically assign new users to groups when you add them to MashZone NextGen. This can simplify some of the process of granting permissions to users.

The feature is not available if you are using your LDAP Directory as the MashZone NextGen User Repository.

You can add both built-in MashZone NextGen groups or groups that you have added for your organization to the list of default groups to automatically assign to new users. If you make **Presto_Developer** a default group, for example, every new user would be granted access to MashZone NextGen dashboard and data feed editors and API console.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Users and Groups** tab and click **Default Groups**.
3. Enter part of a group name and click **Search**.
A list of groups that start with that name displays.
4. Drag a group and drop it in the Default User Groups bucket.
Continue to find and add groups as needed.
5. Click **Save these changes**.

New user will automatically be assigned to all groups in the default list. This change does **not** affect any existing users.

8.3.8.4 Grant dashboard and data feed permissions via API console

You can use the API console to grant the permissions to create and edit dashboards and data feeds to MashZone NextGen users and user groups.

You have two options to assign the permissions to specific users and user groups. Using the Admin Console, as described in Manage Users (page 133) and Manage User Groups (page 134), and in addition, by using the API console.

By default, the permissions to create and edit dashboard and data feeds is only assigned to members of the **Presto_Developer** group, see Built-In MashZone NextGen MashZone NextGen User Groups and Permissions (page 105) for details. By using the API console you can assign the permissions required to other MashZone NextGen users.

Procedure

1. Open the API Console. (page 438)
2. Enter the following code blocks if you want to assign the permissions to a specific user. Enter the relevant user name for the userID parameter.

```
{ "version": "1.1", "sid": "PolicyService", "svcVersion": "0.1", "oid":
"addPermissions", "params": [ "Mashzone", "type.entity.uiObject",
"accessDashboardEditor", [ { "principalId": "userID", "principalTypeId": "User"
} ] ] } { "version": "1.1", "sid": "PolicyService", "svcVersion": "0.1", "oid":
"addPermissions", "params": [ "Mashzone", "type.entity.uiObject",
"accessFeedEditor", [ { "principalId": "userID", "principalTypeId": "User" } ]
] }
```

3. Enter the following code blocks if you want to assign the permissions to a specific user group. Enter the relevant group name for the groupID parameter.

```
{ "version": "1.1", "sid": "PolicyService", "svcVersion": "0.1", "oid":
"addPermissions", "params": [ "Mashzone", "type.entity.uiObject",
"accessDashboardEditor", [ { "principalId": "groupID", "principalTypeId":
"Group" } ] ] } { "version": "1.1", "sid": "PolicyService", "svcVersion": "0.1",
"oid": "addPermissions", "params": [ "Mashzone", "type.entity.uiObject",
"accessFeedEditor", [ { "principalId": "groupID", "principalTypeId": "Group" }
] ] }
```

4. Click **Run**.

The permissions to create and edit dashboards and data feeds are assigned to the relevant user or user group.

8.3.8.4.1 Enable dashboard and data feed creation

The creation of dashboards and data feeds is restricted in MashZone NextGen by default. You can enable a user to create dashboards and data feeds by assigning the user to the **Presto_Developer** group.

If an existing MashZone NextGen repository is used (for example, after a migration to a newer MashZone NextGen version) the permissions might be missing in the **Presto_Developer** group. You can add the required permissions to the repository.

Procedure

1. Open the **PolicyEngineConfiguration.xml** file located in the following directory using a text editor.

```
<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\WEB-INF\classes\
```

2. Add the following code block to the Presto_Developer_PermissionSet and save the file.

```
<Permission resourceName="Mashzone" resourceType="type.entity.uiObject">  
<Action name="accessDashboardEditor"/> <Action name="accessFeedEditor"/>  
</Permission>
```

3. Open the API Console.

4. Enter the following code block.

```
{ "sid": "PolicyService", "oid": "reloadPermissionSets", "version": "1.1",  
  "params": {} }
```

5. Click **Run**.

The permissions are added to the MashZone NextGen repository.

8.3.8.5 Install MashZone NextGen and MashZone NextGen Event Service as Windows services

You can install MashZone NextGen Server and MashZone NextGen Event Service as Windows services by running the **mashzonenextgen-service.bat** file.

The file is located in the <MashZone NextGen installation>\apache-tomcat\bin directory.

Procedure

1. Run **mashzonenextgen-service.bat -install** to install MashZone NextGen Server and MashZone NextGen Event Service as Windows services.
2. Run **mashzonenextgen-service.bat -remove** to uninstall MashZone NextGen Server and MashZone NextGen Event Service as Windows services.

The Windows services **Software AG MashZone NextGen <version>**, **Software AG MashZone NextGen Event Service Core <version>**, **Software AG MashZone NextGen Event Service Zookeeper <version>**, and **Software AG MashZone NextGen Event Service HA-Store <version>** are installed or uninstalled.

8.3.9 Command Central

Using the Software AG Command Central web user interface and command line interface you can manage some MashZone NextGen functionalities in your Software AG products environment. The Command Central is a tool to administer hundreds of managed product installations in your IT

landscape from a central location. Command Central is built on top of Software AG Common Platform and product-specific features are in the form of plug-ins.

MashZone NextGen provides an own Command Central Platform Manager plug-in. The MashZone NextGen plug-in (page 138) is available as a separate item in the SAG Installer product tree. It is automatically selected when installing the MashZone NextGen product.

With MashZone NextGen 10.2, a separate plug-in for MashZone NextGen Event Service (page 141) is available. It is automatically installed with the MashZone NextGen Command Central plug-in.

Detailed information on how to use the Software AG Command Central can be found in the Command Central User Guide.

8.3.9.1 MashZone NextGen plug-in

8.3.9.1.1 Required installer modules

If you want to use the MashZone NextGen Command Central integration, you need to install the following modules from the installer's product tree.

- MashZone NextGen#Server
- Infrastructure#Command Central
- Infrastructure#Platform Manager
- Infrastructure#Platform Manager Plug-Ins#MashZone NextGen Plug-in

8.3.9.1.2 Instance Overview

- The **Instances** tab provides information about the run-time status and installation details.
- By clicking MashZone NextGen on the **Instances** tab you can open the **Overview** page of the MashZone NextGen instance. The overview tab contains information about the runtime status and allows triggering some life-cycle actions like starting and stopping the instance.
- By clicking on the status icon or the link below you can start or stop the MashZone NextGen instance.
- By clicking the flag icon you can open the alerts. You can also clear the alerts by clicking on the number in the circle.

The KPIs shown are available only if the server is configured for remote JMX connections, see chapter **Remote JMX connection** for a configuration example. The following table gives an overview of the supported KPIs.

KPI	Description	Max value	Marginal value	Critical value
Running Calculation	The number of currently running feed calculations	10	4	8

KPI	Description	Max value	Marginal value	Critical value
Heap Memory	The actually used heap memory of the JVM	Configured max. value (parameter "wrapper.java.maxmemory" in wrapper.conf)	80% of max	90% of max
Calc. Queue Size	The number of waiting feed calculations in the queue	20	8	16

8.3.9.1.3 Instance Configuration

The instance configuration supports four different configuration types:

- General properties contained files "presto.config" and "mashzone.properties"
- Java service wrapper properties contained in file "wrapper.conf" resp. "custom_wrapper.conf"
- Java system properties contained in file "wrapper.conf" resp. "custom_wrapper.conf"
- Memory configuration contained in file "wrapper.conf" resp. "custom_wrapper.conf"

The following table gives an overview of the supported configuration types, their location and which configuration aspects they cover.

Configuration type	Configuration instance	Location	Covered aspects
General Properties	MashZone NextGenMashZone NextGen properties	<TOMCAT_ROOT >/webapps/ WEB-INF/mashzone.properties	<ul style="list-style-type: none"> ▪ Initial resource folder ▪ Import/Export folder ▪ JDBC drivers folder ▪ Feed calculation settings ▪ Feed related and general cache settings ▪ Event Service endpoint URL ▪ PPM connection timeout
General Properties	Presto Config	<TOMCAT_ROOT >/webapps/ WEB-INF/classes /presto.config	<ul style="list-style-type: none"> ▪ UI mode "MashZone" or "Classic" ▪ SAML2 configuration ▪ JWT configuration ▪ Landing page welcome text ▪ Parts of the SSL configuration (selfsigned, anyhosts) ▪ Parameters to control DAO/DB code

Configuration type	Configuration instance	Location	Covered aspects
			<ul style="list-style-type: none"> generation Database DAO service registration options
Memory	-	<TOMCAT_ROOT>/conf/wrapper.conf <TOMCAT_ROOT>/conf/custom_wrapper.conf	<ul style="list-style-type: none"> JVM initial heap size JVM maximum heap size
Jaca Service Wrapper	-	<TOMCAT_ROOT>/conf/wrapper.conf <TOMCAT_ROOT>/conf/custom_wrapper.conf	<ul style="list-style-type: none"> Java home and Java command locations Java Main class to be used Java classpath Wrapper log settings Wrapper Display Name Windows service settings Wrapper files locations ...and some more settings of lower importance
Java System Properties	-	<TOMCAT_ROOT>/conf/wrapper.conf <TOMCAT_ROOT>/conf/custom_wrapper.conf	<ul style="list-style-type: none"> Tomcat location settings File encoding Some RAQL settings Presto classic resource folder Terracotta license key and config timeout Derby locks settings Some file locations concerning event services Optional: Remote JMX settings Optional: Remote debugging settings

<TOMCAT_ROOT> usually is <INSTALL_ROOT>/MashZoneNG/apache-tomcat.

The latter three configuration types are handled by the Tanuki Service Wrapper. It supports a "cascaded" configuration. That means, that only the default configuration is contained in file **wrapper.conf**. Parameters that were added or manipulated are stored in file **custom_wrapper.conf** which is included by **wrapper.conf**. If parameters are contained in both files, the ones in **custom_wrapper.conf** overwrite the original values in **wrapper.conf**.

Note: Some important configuration aspects of MashZone NextGen are currently not supported by Command Central. They can be configured by using the MashZone NextGen user interface or by editing specific configuration files. Detailed **information** on how to configure MashZone NextGen can be found in MashZone NextGen Server Configuration (page 148).

8.3.9.1.4 Instance Logs

The **Logs** tab provides the MashZone NextGen's log files located in the `<TOMCAT_HOME>/logs` directory.

8.3.9.1.5 Enable remote JMX connection

In order to enable the remote JMX connection that allows monitoring the KPIs, please add the following parameters via Command Central Configuration Type "Java System Properties" and restart the server using the Command Central life-cycle actions.

```
com.sun.management.jmxremote.port=9607
com.sun.management.jmxremote.authenticate=false
com.sun.management.jmxremote.ssl=false
```

The SPM plug-in detects the JMX port automatically.

8.3.9.2 MashZone NextGen Event Service plug-in

8.3.9.2.1 Instance overview

The Instances tab provides information about the run-time status and installation details.

By clicking **MashZone NextGen Event Service** on the **Instances** tab you can open the **Overview** tab of the MashZone NextGen Event Service instance. The **Overview** tab contains information about the runtime status and allows triggering some life-cycle actions like starting and stopping the instance. By clicking on the status icon or the link below you can start or stop the MashZone NextGen Event Service instance.

By clicking the flag icon you can open the alerts. You can also clear the alerts by clicking on the number in the circle.

8.3.9.2.2 Instance configuration

The instance configuration supports five different configuration types:

- General properties contained in files **rtbs.config**, **zookeeper.properties**, **server.properties**
- Keystore Configuration (**rtbs.properties**)
- Truststore Configuration (**rtbs.properties**)

- Ports (**rtbs.properties**, **zookeeper.properties**, **server.properties**)
- Other: Cluster Configuration (**rtbs.properties**, **zookeeper.properties**, **server.properties**)

The following table gives an overview of the supported configuration types, their location, and which configuration aspects they cover.

Configuration type	Configuration instance	Location	Description
General Properties	Event Service Core Properties	<EventService_ROOT>/conf/rtbs.properties	<p>EDA Configuration</p> <ul style="list-style-type: none"> ▪ EventTypeStore location (com.softwareag.event.routing.eventtypestore.location) <p>DES Configuration</p> <ul style="list-style-type: none"> ▪ TypeRepository location (com.softwareag.events.repository.location) ▪ License location (com.softwareag.events.license.location) ▪ Routing Configuration (com.softwareag.events.routing.configuration.directory) <p>It is not recommended to make other changes to the configuration file. All other settings in the configuration file can be configured using the configuration types listed in this table.</p>
General Properties	HA-Store	<EventService_ROOT>/ha-store/conf/server.properties	<p>Persistence settings for high availability</p> <ul style="list-style-type: none"> ▪ Retention time (log.retention.hours, time in hours before buffered data becomes eligible for deletion). Defaults to 168 hours (7 days) ▪ Retention size (log.retention.bytes, maximum size of buffered data, per buffer). Defaults to infinity, i.e., the size is only bound by retention time) <p>It is not recommended to make other changes to the configuration file. Especially broker.id should only be changed using the Other configuration type (cluster configuration, see below). Port settings should be changed using the ports configuration type (see below).</p>

Configuration type	Configuration instance	Location	Description
General Properties	Zookeeper		Apache Zookeeper configuration. It is not recommended to make any manual changes to the file. All relevant aspects can be configured using the configuration types listed in this table.
Keystores		<EventService_ROOT>/conf/rtbs.properties	Keystore location, alias, type, location, password
Other			<p>Cluster configuration:</p> <ul style="list-style-type: none"> Replication Factor: determines how many replicas of persisted data (for high availability) will be created. The number must be ≥ 1 and \leq the number of Event Service instances in the cluster. The number must be ≥ 2 to achieve any fault tolerance. As a rule of thumb: the higher the number of replicas the more Event Service node failures can be tolerated. However, the amount of storage required on the Event Service machines increases linearly with the number of replicas. <p>Note: The replication factor setting only affects buffers that are created after the replication factor was changed, that is, changing the replication factor does not affect existing buffers.</p> <ul style="list-style-type: none"> Zookeeper-URLs Format: hostname:port, for example, machine1.mycompany.com: 12181 To form a cluster, Zookeeper URLs must contain a comma-separated list of all machines that should be part of the cluster. To form a cluster of the machines machine1, machine2, and machine3, enter: machine1: 12181,machine2: 12181, machine3: 12181 Instance ID The instance ID must be a number between 1 and 255 and must be unique among all cluster members. It is used to identify

Configuration type	Configuration instance	Location	Description
			<p>cluster members. The ID is internally used to configure Event Service Core's rtbs.id, Zookeeper's ID (myid), and HA-Store's broker.id.</p> <ul style="list-style-type: none"> ▪
Ports			Port configuration for Event Service (non-TLS/TLS), Zookeeper, HA-Store
Truststores		<EventService_ROOT>/conf/rtbs.properties	Truststore location, alias, type, location, password, and period. The truststore is required to validate (JWT) tokens that MashZone NextGenMashZone NextGen uses to authenticate against MashZone NextGen Event Service. Make sure to use a truststore that contains a public key matching the private key used in MashZone NextGen (see section Protect MashZone NextGen Event Service access (page 106)).

<EventService_ROOT> usually is <INSTALL_ROOT>/MashZoneNG/rtbs

8.3.9.2.3 Instance Logs

The **Logs** tab provides the MashZone NextGen Event Service log files located in the <EventService_ROOT>/log directory.

8.3.9.3 Configure MashZone NextGen Digital Event Service (DES)

You can configure the URL that subscribes to an external Universal Messaging (UM) instance in MashZone NextGen using the Command Central.

The configuration type **Messaging Services** contains configuration options for Universal Messaging. It is pre-configured to use a Universal-Messaging Server running on local host (URL: nsp://localhost:9000). You must reconfigure the URL if you call a UM instance that is either remote (not localhost) or if the instance is running on another port than the UM default (9000).

Procedure

1. Open a MashZone NextGen instance in Command Central.
2. Open the MashZone NextGen - DES component, a subcomponent of the MashZone NextGen component.

3. Open the **Configuration** tab.
4. Select **Messaging Services** in the drop-down menu.
5. Click **UniversalMessaging** in the **Service Name** column to open the settings.
6. Click **Edit** on the upper right-hand side to edit the settings.
7. Make your settings. For details on the UM settings see Create DES Event Source (Seite 244).
8. Save your settings.
9. Restart the MashZone NextGen server.

Your settings are applied.

If you do not use the Command Central, you can edit the UM URL in a corresponding configuration file stored in the MashZone NextGen installation. Edit the file in a text editor, and replace the value of the messaging provider URL field in the JSON content.

The configuration file is located in the following directory.

<MashZone NextGen

installation>\apache-tomcat\conf\DigitalEventServices\services\UniversalMessaging\service-UniversalMessaging.json

8.3.9.4 Configure ports

You can configure the ports in MashZone NextGen using the Command Central.

By default, the **HTTP** port is enabled, and the **HTTPS** port is disabled.

Procedure

1. Open a MashZone NextGen instance in Command Central.
2. Open the **Configuration** tab.
3. Select **Ports** in the drop-down menu.
4. Click an alias in the **Alias** column to open the port settings.
5. Click **Edit** on the upper right-hand side to edit the port settings.
6. Make your settings. The alias is not editable.
 - a. Additionally, enable the **HTTPS** port to run MashZone NextGen in **HTTPS** mode.
 - b. Deactivate the **HTTP** port and activate the **HTTPS** port to run MashZone NextGen only in **HTTPS** mode.
 - c. Change the port number if required.
 - d. Activating the **HTTPS** port activates the **SSL** mode.

If you want to run the **SSL** mode, you must also configure the Keystores and Truststores in Command Central (page 146).
7. Save your settings.
8. Restart the MashZone NextGen server.

Your settings are applied.

8.3.9.5 Configure keystores

You can configure keystores in MashZone NextGen using the Command Central.

The default keystore type is JKS, but you can change it to PKCS12.

You can only configure the keystores if you previously have activated the SSL mode in the ports configuration (page 145).

Procedure

1. Open a MashZone NextGen instance in Command Central.
2. Open the **Configuration** tab.
3. Select **Keystores** in the drop-down menu.
4. Click an alias in the **Alias** column to open the keystore settings.
5. Click **Edit** on the upper right-hand side to edit the keystore settings.
6. Make your settings. The alias and description are not editable.
 - a. Select the keystore type in the **Type** drop-down menu.
 - b. Specify the location of the keystore file in the file system.
 - c. Specify the password of the keystore, and confirm it.
 - d. Specify a key alias of the key within the keystore you want to use.

If there are multiple keys in the keystore, and you do not specify a key alias, the first one present in the keystore is used.

7. Save your settings.
8. Restart the MashZone NextGen server.

Your settings are applied.

8.3.9.6 Configure truststores

You can configure truststores in MashZone NextGen using the Command Central.

You can only configure the truststores if you previously have activated the SSL mode in the ports configuration (page 145).

Procedure

1. Open a MashZone NextGen instance in Command Central.
2. Open the **Configuration** tab.
3. Select **Truststores** in the drop-down menu.
4. Click an alias in the **Alias** column to open the truststore settings.
5. Click **Edit** on the upper right-hand side to edit the truststore settings.
6. Make your settings. The alias and description are not editable.
 - a. Select the keystore type in the **Type** drop-down menu.
JKS or PKCS12 are currently supported.

- b. Specify the location of the truststore file in the file system.
 - c. Specify the password of the truststore, and confirm it.
7. Save your settings.
 8. Restart the MashZone NextGen server.

Your settings are applied.

8.3.9.7 Configure SSL certificates

You can configure SSL certificates using the Command Central.

Procedure

1. Enable in SSL mode in the port configuration (page 145).
2. Configure a keystore (page 146).
3. Configure a truststore (page 146).

Your settings are applied.

8.3.9.8 Manage licenses

You can manage the licenses of MashZone NextGen and the optional internally used product Terracotta (BigMemory) using the Command Central.

The list contains always two entries, one for the MashZone NextGen license and one for the optional internal Terracotta license.

- The first column shows the license type, either MashZone NextGen or Terracotta.
- The second column shows the current license status. This is either **Not present** if no license is available, **Valid** if the license is valid, **Invalid** if the license is invalid, **Expiring soon** if the license will expire within the next few days.
- The third column shows the actual expiration date of the license. If a license has no expiration date, this field will remain empty.

Procedure

1. Open a MashZone NextGen instance in Command Central.
2. Open the **Configuration** tab.
3. Select **Licenses** in the drop-down menu.
4. Click a license type in the **License Type** column to open the license settings.
5. Click **Edit** on the upper right-hand side to edit the license settings.

A drop-down box is now displayed that shows the currently available licenses in the Command Central license repository. The license list contains only the licenses for the license type you have selected. If no license is offered, you must first upload a license to the

Command Central license repository. Detailed information on how to upload a license file can be found in the Command Central User Guide.

6. Select a license in the drop-down list.
7. Save your settings.
8. Restart the MashZone NextGen server.

Your settings are applied.

8.4 MashZone NextGen Server Configuration

8.4.1 Memory Configuration for the MashZone NextGen Server

MashZone NextGen is initially installed with default memory settings for a small web application. Your actual memory requirements may vary significantly based on your expected load, throughput and environment.

Note: With release 3.7, MashZone NextGen is no longer supported on 32-bit architectures which have memory access limitations from Java.

If you are working with large datasets in MashZone NextGen Analytics or you are deploying MashZone NextGen in a staging or production environment, you may need to tune this default memory configuration. Which options you can configure depends on your usage and environment considerations:

If	See
<ul style="list-style-type: none"> ▪ Your computer has less than 4G of RAM memory, or ▪ You have not installed BigMemory Server(s). 	Configuration When MashZone NextGen Uses Only Heap Memory (page 148)
<ul style="list-style-type: none"> ▪ You have installed BigMemory Server(s), or ▪ Your computer has more than 4G of RAM memory. 	Configuration When MashZone NextGen Uses Heap and Off-Heap Memory (page 149)

8.4.1.1 Configuration When MashZone NextGen Uses Only Heap Memory

The initial default Java heap memory settings for MashZone NextGen are appropriate for small applications or development environments 2G as the maximum heap size. The Event Service that is deployed with MashZone NextGen also uses this heap space.

If you are using MashZone NextGen with large datasets and limited memory (less than the recommended 4G minimum), configuring BigMemory to use off-heap memory is **not** recommended as it can adversely affect performance.

Some portion of available memory must be reserved for the operating system and any other applications on this host.

How much memory to allocate to MashZone NextGen depends on available memory, requirement for the Event Service and what other applications may run on this computer.

Procedure

1. In a text editor of your choice, open the application server configuration file `MashZoneNG-install/apache-tomcat/conf/wrapper.conf`
2. Change any of these Java memory options:
`wrapper.java.initmemory`, Default = **512**
`wrapper.java.maxmemory`, Default = **2048**
3. Save your changes and restart the MashZone NextGen Server. See [Start and Stop the MashZone NextGen Server](#) (page 113) for instructions.

8.4.1.2 Configuration When MashZone NextGen Uses Heap and Off-Heap Memory

MashZone NextGen should be configured to use both heap and off-heap memory only when the available memory supports this adequately and you have also installed BigMemory Servers.

You must have installed a copy of your BigMemory license in MashZone NextGen to use off-heap memory. See [Manage Licenses for MashZone NextGen and BigMemory](#) (page 115) for instructions.

With combination heap and off-heap memory, as this figure shows, BigMemory uses off-heap memory for the MashZone NextGen Analytics In-Memory Stores and MashZone NextGen MashZone NextGen caches. All other MashZone NextGen processing, including the Event Service that is deployed with MashZone NextGen, remains in heap.

The total available off-heap memory may be limited to local off-heap memory as shown above, or it may include additional off-heap memory on external hosts if you have installed BigMemory Server arrays.

Procedure

1. In a text editor of your choice, open the application server configuration file `MashZoneNG-install/apache-tomcat/conf/wrapper.conf`.
2. Change or add either of these memory options used with BigMemory, see table below.
3. Change or set any of these Java memory options.
`wrapper.java.initmemory`, Default = **512M**
`wrapper.java.maxmemory`, Default = **2G**
See the [Java Tuning White Paper](#) for more information and suggestions.
4. Save your changes and restart the MashZone NextGen Server. See [Start and Stop the MashZone NextGen Server](#) (page 113) for instructions.

Memory options used with BigMemory	Description
wrapper.java.additional.<n+1> =-Dpresto.bm.maxOffHeap	Default = 1G Where n is the number of last additional Java parameter. This is the maximum size of local off-heap memory that BigMemory can use for the MashZone NextGen Analytics In-Memory Stores and MashZone NextGen caches. This property sets off-heap memory limits in the MashZoneNG-config/ehcache.xml configuration file. The total size of off-heap memory may include additional, external memory depending on how BigMemory is deployed.
wrapper.java.additional.<n+2> =-XX:MaxDirectMemorySize	Default = 1500M Where n is the number of last additional Java parameter. This Java memory option must be set to allow access to both off-heap and an additional allocation for Java. The value of this option must always be larger than the memory allocated to off-heap. A good rule of thumb is at least 500M more.

8.4.2 Support International Character Sets and Locales

International character sets are the alphabets, characters, glyphs and other symbols used in any non-English language. Technically, this includes any non-ASCII characters. To handle these characters properly, MashZone NextGen must know the character set and how it is digitally represented - the **character encoding**.

MashZone NextGen uses the UTF-8 character encoding to handle character sets for all languages. Both UTF-8 and UTF-16 can represent any Unicode character.

Unicode defines a unique encoding for every character in world languages that are currently in active use as well as some well-known dead languages, such as ancient Greek.

Locale identifies the language used and potentially specific regional spelling or usage aspects of the language, such as differences between American English (EN_us) versus Australian English (EN_au). Locale also identifies the formats used to present dates, times and numbers for that region.

Both the character encoding and locale help to ensure that MashZone NextGen properly handles and presents data to users. The areas of configuration involved in this support include:

- MashZone NextGen Repository: the character encoding for this repository determines what character sets users can use when they create artifacts. The timezone for the MashZone NextGen Repository also affects timestamps shown in MashZone NextGen.

- **Display options:** in most cases, date, time and numeric data are shown to users based on browser settings or a default locale. Some views allow users to choose date and time formats.
- See Date, Time and Numeric Display Options (page 152) for details.
- **Logging:** you can also support international characters and different locales for the messages sent to MashZone NextGen logs. See Message Log and Default Locales (page 152) for details.

8.4.2.1 Set the Repository Character Encoding

The character encoding for the MashZone NextGen Repository is defined when you create the database that will host the repository. To support international character sets, this should be set to UTF-8, or for some databases UTF-16.

Important: Because of known issues, artifact names and the IDs that are generated from these names are restricted to ASCII characters. The syntax and encoding names you must use are specific to each database. For more information, see:

- Documentation for your database
- And either:
 - Move MashZone NextGen repository to Microsoft SQL Server (page 118)
 - Move the MashZone NextGen repository to MySQL (page 120)
 - Move the MashZone NextGen repository to Oracle (page 122)
 - Move the MashZone NextGen repository to PostGres (page 124)

8.4.2.2 Set the Repository Timezone or Offset

The default timezone that is used to record timestamps such as the created date and time for an artifact is the timezone for the host of the MashZone NextGen Repository. You can change the timezone used to save repository timestamps or set an offset so that repository timestamps are displayed in a different timezone:

- Force the timezone that the MashZone NextGen Repository uses to match the timezone for the MashZone NextGen Server. See Synchronize the MashZone NextGen Repository and MashZone NextGen MashZone NextGen Server Time Zones (page 268) for instructions.
- Configure the display timezone in MashZone NextGen as an offset from UTC. Note that this does not affect the actual timezone recorded in the MashZone NextGen Repository.
- This offset is defined in the **repositoryTimezoneOffset** property in `web-apps-home/mashzone/hub/config.js` file. This property is undefined by default.
- Edit this JSON property, setting the number of minutes as a UTC offset. For example, **300** sets this to Eastern Standard Time while **-180** sets this to Arabic Standard Time.
- Once the property is saved, restart the MashZone NextGen Server.

8.4.2.3 Date, Time and Numeric Display Options

In general, MashZone NextGen displays dates, times and numeric data using the formats defined by the browser's locale for the current user. If no locale information is available from the browser, MashZone NextGen use the default system locale which typically is EN_us.

Some built-in views in MashZone NextGen, allow users to choose a date and time pattern for result data such as **mm/yyyy**, for the month and year, or **EEE MMM dd, yyyy**, for the day of the week, month name, day and year. This pattern determines the widgets of dates or times that display in that view, but the language, order and delimiters used in the display are determined by the user's locale or the default locale.

8.4.2.4 Message Log and Default Locales

MashZone NextGen uses the default locale defined for the JVM for all messages that are added to MashZone NextGen logs. These defaults may also be used as the locale for artifacts if no locale is defined by users or provided by the client browser.

To set the locale for the JVM for MashZone NextGen using Java properties

- In a text editor of your choice, open the file
<MashZoneNG-installation>/apache-tomcat/conf/wrapper.conf.
- Add the lines just below the line which sets the last additional Java parameter:
- `wrapper.java.additional.<n+1>=-Duser.country=country-code`
- `wrapper.java.additional.<n+2>=-Duser.language=language-code`
- Where n is the number of last additional Java parameter.

For example: **wrapper.java.additional.20=-Duser.country=CA**
wrapper.java.additional.21=-Duser.language=fr

- Save your changes to the file.
- Restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.4.3 Edit style templates

You have MashZone NextGen administrator permissions.

You can edit the style templates supplied with MashZone NextGen. Editing the style templates enables you to customize the look and feel of your dashboards and the dashboard editor, for example, colors schemes, fonts, brand logo, or background colors.

If you change the style template at the application level your settings are also applied in the MashZone NextGen Feed Editor.

The application style template file **application.less** is located in the following folder on the MashZone NextGen server.

<MashZone NextGen

installation>\apache-tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\application.

The dashboard style template files are located in the following folder on the MashZone NextGen server. The default dashboard template file is named **default.less**.

<MashZone NextGen

installation>\apache-tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\dashboard

The style templates provide basic variables that are used as default variables by many widgets, for example, @header-color. You can override the basic variables by defining more specific, customized variables, for example, @header-dashboard-name-color.

For example, if @header-color is black and @header-dashboard-name-color is red, and both variables are the only ones defined, all labels in the header will be black, except for the dashboard name label, which will be red.

All application and dashboard variables available are described in the corresponding style template files.

Procedure

1. Open the relevant template style file in your text editor.
2. Edit the style parameters according your requirements.
3. Save your changes.
4. Reload the changed style template files.
 - a. Open a dashboard in the MashZone NextGen Dashboard Editor.
 - b. Click **Manage > Change style template** in the dashboard main menu.
 - c. To reload the application style template click **Activate**.
 - d. To reload the dashboard style template select the relevant dashboard style template in the drop-down menu and click **Update**.
5. Click **OK**.

The changed style template has been reloaded and the style templates are applied to the application or dashboard.

8.4.4 Configure the MashZone NextGen server with custom ports

Port configuration is initially set when you install MashZone NextGen. If you change these ports or you need to host multiple MashZone NextGen Servers on one host, you must update configuration in the MashZone NextGen Server. You may also need to change ports for the MashZone NextGen Repository and for Tomcat.

8.4.4.1 Change MashZone NextGen Server Ports

The host name and port for the MashZone NextGen Server defaults to localhost and 8080 respectively. The port is typically defined in configuration for Tomcat, the application server that hosts MashZone NextGen.

Procedure

1. Update port configuration for the application server that hosts MashZone NextGen in MashZoneNG-install/apache-tomcat/conf/server.xml in the <Connector> elements for HTTP and/or HTTPS.
2. Restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.4.4.2 Change MashZone NextGen Repository Ports

If you are running multiple instances of the MashZone NextGen Server in one host with separate instances of the MashZone NextGen Repository, you may need to use different ports for each database instance.

- The default MashZone NextGen Repository. No updates to ports are needed as the Derby database is embedded.
- The MashZone NextGen Repository is hosted in a robust database. You must use different ports for each MashZone NextGen Repository instance.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **MashZone NextGen Repositories** section.
3. Select **Metadata Repository** and:
 - a. Change the JDBC URL to the correct host and port number.
 - b. Click **Save**.
4. Restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.4.4.3 Tomcat Application Server Port

If you are running multiple instances of the MashZone NextGen Server in one host, you must have separate application server instances for each. You must update the following ports:

Configuration File	Element
MashZoneNG-install/apache-tomcat/conf/server.xml	< Server > to set the command port < Connector > for the HTTP port < Connector > for the HTTPS port

8.4.5 Configure the MashZone NextGen server to work with a proxy server

MashZone NextGen is **only** compatible with HTTP proxy servers.

If you have a proxy server in your environment that MashZone NextGen should use, you **must** add configuration information to the MashZone NextGen Server for the proxy server. You can also define a **whitelist** of addresses that do not require proxy server access. See Define a Proxy Server Whitelist for MashZone NextGen (page 157) for more information.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. In the Server section, click **Proxy Settings**.
3. Set the **Enable Proxy** option.
4. Set the following connection properties for your proxy server:
 - Host** = the host name or IP address for the proxy server. This is required.
 - Port** = the port number for the proxy server. This is required.
 - Username** = the user name that the MashZone NextGen Server should use to connect to the proxy server. This is only required if your proxy server requires credentials.
 - Password** = the password that the MashZone NextGen Server should use to connect to the proxy server. This is only required if your proxy server requires credentials.
5. If needed, define a whitelist of addresses that should not use the proxy server. See Define a Proxy Server Whitelist for MashZone NextGen (page 157) for instructions.
6. Click Save proxy settings.

8.4.6 Embedding MashZone NextGen in external system environments

You can use MashZone NextGen as a widget in external products, for example, webMethods Business Console. As embedded widget MashZone NextGen is enabled to send data via outbound API (Post data) to the embedding system and receive data via inbound API (URL selection) from the embedding system.

8.4.6.1 Configure MashZone NextGen server to work with iFrame

By default, MashZone NextGen can be embedded using HTML inline frames (iFrame) if the MashZone NextGen server and the server of the embedding system use the same protocol, same host and same port.

To embed MashZone NextGen within another HTML document the iFrame source points to the MashZone NextGen dashboard as shown in the following example. `<iframe`

```
id="embedded-mzng-dashboard" width="600px" height="600px"
src="http://mzngServerHost:mzngServerPort/mashzone/hub/
dashboard/dashboard.jsp?mzngDashboardGUID"> <p>Your browser does not
support iframes.</p> </iframe>
```

If the embedding system is running on a different host or uses a different protocol or port, the MashZone NextGen server must be configured as follows. The MashZone NextGen MashZone NextGen server configuration file **applicationContext-security-filters.xml** needs to be configured by adding filters for **X-Frame-Options** and content security policies.

The **applicationContext-security-filters.xml** server configuration file is located in following directory. `<MashZone NextGenMashZone NextGen-install>/apache-tomcat/webapps/mashzone/WEB-INF/classes.`

Procedure

1. Open the `applicationContext-security-filters.xml` configuration file in a text editor of your choice.
2. Adapt the security settings as follows and exchange the string `"http://otherServerHost:otherServerPort"` with the system origin MashZone NextGen is to be embedded in.

```
<beans:beans xmlns="http://www.springframework.org/schema/security"...> ...
<http pattern="/hub/(login|reset_password)\.html.*" security="none"
request-matcher="regex"/> <http pattern="/help/*.html.*" security="none"
request-matcher="regex"/> <http pattern="/**/*.*.jsp" use-expressions="false"
authentication-manager-ref="authenticationManager"
entry-point-ref="mzngAuthenticationEntryPoint"> <anonymous enabled="false"/>
<headers> <!--frame-options policy="SAMEORIGIN"/--> <frame-options
policy="ALLOW-FROM" strategy="static"
value="http://otherServerHost:otherServerPort" /> <!--content-security-policy
policy-directives="frame-ancestors 'self'"/--> <content-security-policy
policy-directives="frame-ancestors 'self'
http://otherServerHost:otherServerPort"/> </headers> <csrf
token-repository-ref="csrfTokenRepository"
request-matcher-ref="skipHttpAuthCsrfMatcher"/> <custom-filter
ref="samlTokenProcessingFilter" after="PRE_AUTH_FILTER"/> <custom-filter
ref="jwtTokenProcessingFilter" before="CAS_FILTER"/> <custom-filter
ref="credentialContainerFilter" before="EXCEPTION_TRANSLATION_FILTER"/>
</http> <http pattern="/**/*.*.html" use-expressions="false"
authentication-manager-ref="authenticationManager"
entry-point-ref="mzngAuthenticationEntryPoint"> <intercept-url
pattern="/**/*.*.html" access="IS_AUTHENTICATED_ANONYMOUSLY"/> <anonymous
enabled="false"/> <headers> <!--frame-options policy="SAMEORIGIN"/-->
<frame-options policy="ALLOW-FROM" strategy="static"
value="http://otherServerHost:otherServerPort" /> <!--content-security-policy
```

```
policy-directives="frame-ancestors 'self'"/--> <content-security-policy
policy-directives="frame-ancestors 'self'
http://otherServerHost:otherServerPort"/> </headers> </http> ... </beans:beans>
```

3. Save changes.

Your changes will be applied with the next MashZone NextGen server start.

Further details on the topic **Using iFrame** can be found in the spring security documentation: <https://docs.spring.io/spring-security/site/docs/current/reference/html/headers.html#headers-frame-options>.

8.4.6.2 Post data

The **Post data** action creates an outbound API to pass data from MashZone NextGen dashboards to an embedding system, for example, an external web application. The action is available for most widgets.

8.4.6.3 URL selection

With the **URL selection** MashZone NextGen provides an inbound API to receive data from an embedding system, for example, an external web application. The action is available for most widgets.

8.4.7 Define a Proxy Server Whitelist for MashZone NextGen

If you have configured a proxy server for the MashZone NextGen Server, you can define a whitelist of domains, hosts or IP addresses that do **not** require access through the proxy server.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. In the **Server** section, click **Proxy Settings**.
3. Set the following properties:

Bypass IP List = enter one or more IP address, separated by commas, that the MashZone NextGen Server should access without the proxy server. To use wildcards in IP addresses, see Using Regular Expressions in a Whitelist (page 158).

Bypass Host List = enter one or more fully-qualified host names, separated by commas, that the MashZone NextGen Server should access without the proxy server. To use wildcards in IP addresses, see Using Regular Expressions in a Whitelist (page 158).

Bypass Domain List = enter one or more domain names, separated by commas, that the MashZone NextGen Server should access without the proxy server. To use wildcards in IP addresses, see Using Regular Expressions in a Whitelist (page 158).

4. Click **Save proxy settings**.

5. Log out of MashZone NextGen.
6. Stop and restart the MashZone NextGen Server to apply these changes.

8.4.7.1 Using Regular Expressions in a Whitelist

All of the whitelist properties accept **regular expressions** to define sets of IP, host or domain name addresses using wildcards. You can use any valid regular expression character, however, the most common wildcard characters that you may use are:

Replace Any Single Character	.
Replace Any Number of Characters	.*

Examples and solutions for the most common patterns include:

- Specifying Literal Dot Separators (page 158)
- Specifying Domains (page 158)
- Specifying Host Names (page 159)

8.4.7.2 Specifying Literal Dot Separators

Because the dot character is used as a wildcard in regular expressions and is also the standard separator for groups in IP addresses, domain names and host names, you can get unintended results when using wildcards. For example, this is a valid regular expression for IP addresses:

139.16.1.*

On Windows systems, many administrators would expect this to expression to "match any IP address with first-through-third groups of 139, 16 and 1 respectively" such as **139.16.1.10** and **139.16.1.35**.

This would actually match either of these IP addresses:

139.16.1.10

139.16.11.120

In most cases, the difference between a literal dot and the dot as a wildcard character doesn't make a difference. If you need to clarify a whitelist entry to match a literal dot, use **\.** instead.

The expression **139\.16\.1\..*** would correctly match **139.16.1.10** and **139.16.1.35** but would not match **139.16.11.120**. In many cases, you could also simplify this to **139.16.1\..*** to get the correct behavior.

8.4.7.3 Specifying Domains

With domains, you must specify a wildcard at the beginning of the domain name. This example is not a valid domain name expression:

mydomain.com

This entry would **not** match a host name of **east.mydomain.com** or **east.customers.mydomain.com**. To specify the domain correctly, enter:

. *mydomain.com

8.4.7.4 Specifying Host Names

In whitelist properties, host names are fully-qualified. Thus **stives** is not a valid host name while **stives.customers.mydomain.com** is valid. A host name expression of **. *customers.mydomain.com** would match all of these hosts:

stives.customers.mydomain.com

cour.customers.mydomain.com

tempcustomers.mydomain.com

Note that an expression of **. *.customers.mydomain.com** would also match these same three hosts.

You may need to specify literal dot separators in host names also to properly clarify the expressions. If in this example you did **not** want **tempcustomers.mydomain.com** to be matched, you would need an expression such as **. *\ .customers.mydomain.com**. See [Specifying Literal Dot Separators](#) (page 158) for more information.

8.4.8 Configure MashZone NextGen for SSL and Digital Certificates

MashZone NextGen expects HTTP as the default transport protocol from clients to the MashZone NextGen Server.

MashZone NextGen supports HTTPS and SSL for connections from clients or connections to many types of information sources.

The certificate store, certificates and configuration needed to support SSL in MashZone NextGen depends on the connection requirements, as shown below:

	Certificate Store and Certificates		Store Configuration			MashZone NextGenMashZone NextGen Configuration	
	Key	Trust	Java	App Server	MashZone NextGenMashZone NextGen	Authenti cation	Securit y Profiles
One-Way SSL to MashZone NextGen (page 161).	✓			✓			

Mutual to MashZone NextGen See Configure Mutual SSL Between Users and MashZone NextGen (page 161).	✓	✓		✓		✓	
One-Way SSL to Mashable Information Sources (page 161).		✓	can be in either				

See also The Certificate Store and Certificates (page 160) for more information:

8.4.8.1 The Certificate Store and Certificates

Both key stores and trust stores are **certificate stores** to store and manage the **key certificate pairs** or **public certificates** used in secure connections with the SSL protocol. Key stores manage key certificate pairs and trust stores manage the public certificates of trusted peers.

8.4.8.2 Key Certificate Pairs

For MashZone NextGen, the key certificate pair stored in the key store identifies the MashZone NextGen Server to users, for both one-way and mutual SSL. The key certificate pair identifies the MashZone NextGen Server to information sources for mutual SSL.

You must generate a key certificate pair for MashZone NextGen. Typically you also have the key certificate pair signed by a Certificate Authority and import this into the certificate store using the Java **keytool** utility or other certificate management tools.

8.4.8.3 Trusted Peer Certificates

The public certificates from peers are stored in the trust store and identify users, for mutual SSL, or identify information sources, for one-way or mutual SSL.

When public certificates for peers are signed by well known Certificate Authorities, they are automatically verified and imported into the trust store. If public certificates are self-signed or signed by an unknown Certificate Authority (the CA root certificate is not found in the trust store), you must obtain and import the public certificates to the trust store before the first connection occurs during user login.

8.4.8.4 The Certificate Store

You can use a single certificate store as both the key store and trust store for MashZone NextGen or you can use separate certificate stores. You can use an existing certificate store for MashZone NextGen, such as the default certificate store shipped with some application servers. Or you can create a new certificate store using the Java **keytool** utility.

See Java keytool documentation

(<http://docs.oracle.com/javase/6/docs/technotes/tools/solaris/keytool.html>) for more information, commands and instructions on managing key certificate pairs, trusted certificates and certificate stores.

8.4.8.5 Configure Mutual SSL Between Users and MashZone NextGen

The MashZone NextGen Server and users both exchange certificates. MashZone NextGen can also be configured to use user digital certificates for authentication. The connection requires:

- **Store and Certificates:**
 - A certificate store as key store and trust store for the MashZone NextGen Server.
 - A key certificate pair for the MashZone NextGen Server.
 - Public certificates in the trust store for any user public certificates that are self-signed.
 - You must add self-signed certificates to the trust store before these users login. See [Trusted Peer Certificates \(page 160\)](#) for more information.
 - See [The Certificate Store and Certificates \(page 160\)](#) for more information.
- Configuration in the application server hosting MashZone NextGen to use the HTTPS port. This also includes configuration identifying the key store and trust store for the MashZone NextGen Server. See [Configure HTTPS and Certificate Stores in the Application Server \(page 162\)](#) for instructions.

8.4.8.6 One-Way SSL to MashZone NextGen

This requires:

1. A key store and a key certificate pair for MashZone NextGen. See [The Certificate Store and Certificates \(page 160\)](#) for more information.
2. Configuration in your application server for the HTTPS port to MashZone NextGen and the key store. See [Configure HTTPS and Certificate Stores in the Application Server \(page 162\)](#) for instructions.

8.4.8.7 One-Way SSL to Information Sources

This requires:

- A trust store for MashZone NextGen. See [The Certificate Store and Certificates \(page 160\)](#) for more information.
 - Configuration for the trust store in either:
 - The application server hosting the MashZone NextGen Server. See [Configure HTTPS and Certificate Stores in the Application Server \(page 162\)](#) for instructions.

- Java. See Update SSL Configuration for Java (page 163) for instructions.
- Self-signed certificates, if any, for mashable information source using one-way SSL. You must add these certificates to the trust store before the mashable information source can be registered. See Trusted Peer Certificates (page 160) for more information.

8.4.8.8 Configure HTTPS and Certificate Stores in the Application Server

Configuration for SSL for MashZone NextGen can be defined in the application server that hosts the MashZone NextGen Server. These instructions discuss the basic steps for configuring SSL in Tomcat. See Tomcat documentation

(http://tomcat.apache.org/tomcat-8.0-doc/ssl-howto.html#SSL_and_Tomcat) or the documentation for your application server for detailed information.

Procedure

1. If you do not yet have a key store, trust store and certificate for the MashZone NextGen Server, find or create these stores and certificate. See The Certificate Store and Certificates (page 160) for instructions.
2. Configure Tomcat for secure connections from clients to the MashZone NextGen Server.

Warning

We recommend that you do not use TLSv1.0 in your production environments, as TLSv1.0 is considered unsafe for network communication.

You can add the **sslEnabledProtocols** and **sslProtocols** parameters to the **server.xml** file to disable TLSv1.0. The parameters allow limiting the supported protocols. For details, see the Tomcat documentation

(http://tomcat.apache.org/tomcat-8.0-doc/ssl-howto.html#SSL_and_Tomcat).

- a. Edit the **server.xml** file for Tomcat to uncomment and configure the **<Connector>** element for SSL/HTTPS 1.1. For example:

```
<Connector port="8443"
protocol="org.apache.coyote.http11.Http11Nio2Protocol"
SSLEnabled="true" maxHeaderSize="8192"
maxThreads="150" minSpareThreads="25" maxSpareThreads="75"
enableLookups="false" disableUploadTimeout="true"
acceptCount="100" scheme="https" secure="true"
clientAuth="true" sslProtocol="TLS" keystoreFile="conf/tomcat.jks"
keystorePass="keystpwd" truststoreFile="conf/tomcat.jks"
truststorePass="trustpwd" />
```

This example uses the default Tomcat port, 8443, and mutual SSL, based on the **clientAuth** value. If this was a one-way connection, you would set **clientAuth** to false. This example also uses the default Tomcat certificate store, **conf/tomcat.jks**, as both the key store and the trust store. See Tomcat documentation for information on other properties.

- b. Once you have configured an HTTPS port in your application server, update port configuration for the MashZone NextGen Server to listen to that port. See Configure the

MashZone NextGen server with custom ports (page 153) for more information on this step.

- c. Enable MashZone NextGen to use secure session cookies:

Open the **web.xml** file located in **<MashZone NextGen installation>/apache-tomcat/webapps/mashzone/WEB-INF/** in a text editor.

Find the **session-config/cookie-config/secure** element and change the value to **true**.

Example

```
<session-config> <session-timeout>30</session-timeout> <!-- Set the "secure"
flag to true when using HTTPS for enhanced security --> <cookie-config>
<secure>>false</secure> </cookie-config> </session-config>
```

Once this is set to true, only HTTPS access will be allowed.

8.4.8.9 Update SSL Configuration for Java

For the data acquisition using HTTPS, certificates for secure endpoints are validated against the default trust store for Java (JRE).

For more information on the default JRE trust store, see

<http://download.oracle.com/javase/6/docs/technotes/guides/security/jsse/JSSERefGuide.html#CustomizingStores>

(<http://download.oracle.com/javase/6/docs/technotes/guides/security/jsse/JSSERefGuide.html#CustomizingStores>).

Initially, this may not be the trust store you have configured for the MashZone NextGen Server in the application server. This can cause security errors for information sources.

To avoid these errors, you can configure the JRE to use the trust store for the MashZone NextGen Server.

Procedure

1. Open the application server configuration file `MashZoneNG-install/apache-tomcat/conf/wrapper.conf` in a text editor of your choice.
2. Add the following Java system properties:
`wrapper.java.additional.<n+1>=-Djavax.net.ssl.trustStore=/path/to/mashup-server/truststore`
This is the absolute path to the trust store for the MashZone NextGen Server.
`wrapper.java.additional.<n+2>=-Djavax.net.ssl.trustStorePassword=truststore-password`
This is only required if the MashZone NextGen Server's trust store uses a password.
Where n is the number of last additional Java parameter.
3. Save your changes to the script and restart the MashZone NextGen Server. See [Start and Stop the MashZone NextGen Server](#) (page 113) for instructions.

8.4.9 MashZone NextGen Logging

In addition to logging from your application server, MashZone NextGen provides the following types of logging:

- Basic logging for MashZone NextGen Server startup, shutdown and exceptions based on a configured logging level. See [Configure Logging for the MashZone NextGen Server](#) (page 164) for more information.
- Audit logging for dashboards, data feeds, aliases, and permissions. See [Audit logging for dashboards, data feeds, aliases, and permissions](#) (page 165) for details.

8.4.9.1 Configure Logging for the MashZone NextGen Server

The MashZone NextGen Server log, **prestoserver.log**, logs all exceptions from startup through shutdown. See [MashZone NextGen Logging](#) (page 164) for links to additional types of logging you can use with MashZone NextGen.

For clustered environments, updating logging configuration affects logging only for the MashZone NextGen Server where you are currently logged in. Generally, this is the behavior you want.

To change logging for all MashZone NextGen Servers in the cluster, update logging configuration for one server and copy the updated `MashZoneNG-install/apache-tomcat/conf/log4j.properties` file to each of the other MashZone NextGen Servers in the cluster. You do not need to restart MashZone NextGen Servers.

Procedure

1. Open the **log4j** configuration file in `<MashZone NextGen installation>/apache-tomcat/conf/log4j.properties` with a text editor of your choice.
2. Edit any of the following properties:

Root Log Warning = the general logging level to use, such as **ERROR**. All exceptions for that level and above will be logged, so this contributes directly to how quickly logs may grow. **DEBUG** is the most verbose logging level.

Log file path = both the folder where the log files for the MashZone NextGen Server should be saved and the name to use for log files. This defaults to `tomcat-install/logs/prestoserver.log`.

You can use an absolute path or a relative path. Relative paths are relative to the `web-apps-home` folder.

Maximum log file size = maximum size for a log file. Once a file has reached this size it is saved as a numbered backup, such as **prestoserver.log.1** and a new log file is started.

Data nucleus logging level = This property should only be changed when requested by MashZone NextGen technical support to help debug specific issues. It is the logging level to use in the data mapping layer for MashZone NextGen.

HTTP client logging level = This property should only be changed when requested by MashZone NextGen technical support to help debug specific issues. It is the logging level to use for requests/responses between the MashZone NextGen Server and information sources.

NET SF logging level = This property should only be changed when requested by MashZone NextGen technical support to help debug specific issues. It is the logging level to use for JSON serialization and deserialization.

ACEGI security logging level = This property should only be changed when requested by MashZone NextGen technical support to help debug specific issues. It is the logging level to use with the MashZone NextGen security layer.

Apache logging level = This property should only be changed when requested by MashZone NextGen technical support to help debug specific issues. It is the logging level to use with many of the third party libraries used in MashZone NextGen.

Spring framework logging level = This property should only be changed when requested by MashZone NextGen technical support to help debug specific issues. It is the logging level to use for server initialization, shutdown and the request/response pipeline for the MashZone NextGen Server.

3. If desired, click **Advanced Options** to set any of these properties:

Log class for normal logging = the java class that handles appending log entries to the log file. This defaults to **org.apache.log4j.RollingFileAppender**.

Layout class for normal logging = the Java class that handles the layout pattern for entries to log files. This defaults to **org.apache.log4j.PatternLayout**.

Layout pattern for normal logging = the expression defining the pattern for entries to the log file. This defaults to **%d %p [%c - %m%n**

Maximum normal log file backups = how many log backups are kept. This defaults to seven.

The advanced properties are defined by Log4J, the underlying logging framework for MashZone NextGen. For more information, see <http://logging.apache.org/log4j/1.2>.

This change becomes effective automatically within 60 seconds.

8.4.9.2 Audit logging for dashboards, data feeds, aliases, and permissions

The audit logging tracks various events concerning dashboards, data feeds, aliases, and permissions. Logins, logouts and failed logins are also tracked. This logging is enabled by default.

The following events are logged.

- create, edit, and delete of dashboards
- create, edit, and delete of data feeds
- create, edit, and delete of aliases
- edit permissions
- logins, logouts and failed logins

The tracked events are logged in the **MashZone-AuditLog.log** file. The file is located in the following directory.

<MashZone NextGen installation>/logs

The audit logging can be switched on and off in the **log4j2.xml** file. You can edit the XML file with any text or XML editor. If you want to disable the audit logging, set the value of the **level** parameter to **OFF**. The default value is **info**.

The file is located in the following directory.

<MashZone NextGen installation>/conf

8.4.10 MashZone NextGen Notifications

MashZone NextGen sends notifications to users for many events, such as change password requests, password changes, etc.

You can also configure MashZone NextGen to send notifications as email messages.

Procedure

1. Configure an email server in MashZone NextGen. See Configuring a Mail Server for MashZone NextGen (page 166) for instructions.
2. Optionally define the host and port to use for MashZone NextGen in links within email notices. See Configure the MashZone NextGen server with custom ports (page 153) for more information.
3. Turn on email notifications in the Admin Console:
 - a. Open the MashZone NextGen Admin Console (page 95).
 - b. If needed, expand the Server section and click Mail Notifications.
 - c. Set the Turn ON Notifications option and click Save mail notification settings.
4. If you are using your LDAP Directory as the MashZone NextGen User Repository, MashZone NextGen expects to find the email address for a user in the mail attribute in LDAP user entries. If your LDAP Directory stores email addresses in a different LDAP user attribute, you must configure MashZone NextGen to look for the correct LDAP user attribute. See Update the User Email Attribute from LDAP (page 167) for instructions.

8.4.10.1 Configuring a Mail Server for MashZone NextGen

By default, MashZone NextGen sends notices for a variety of administrator and user actions within MashZone NextGen. If you choose to send these notifications as email message, you must add connection configuration to the MashZone NextGen Server for your mail server.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. In the Server section, select Mail Server.
3. Complete the connection properties:

Host = your mail server domain name or host IP address.

Port = the port number for your mail server. This defaults to 25, for SMTP.

4. Set the Requires authentication if your SMTP server requires credentials and complete these properties:

Username = the user account the MashZone NextGen Server should use to connect to your mail server.

Password = the password the MashZone NextGen Server should use to connect to your mail server.

5. If your mail server uses a secure connection, choose the protocol for Connection security:

STARTTLS = a transport layer security (TLS) that does not require a different port for the mail server.

SSL/TLS = transport layer security (TLS) using the secure socket layer (SSL).

6. Click Save mail server settings.

8.4.10.2 Update the User Email Attribute from LDAP

When MashZone NextGen is configured to use your LDAP Directory as the User Repository, MashZone NextGen is configured by default to expect user email addresses in the **mail** attribute for LDAP user entries. If your LDAP Directory uses a different user attribute to store email addresses, you must update LDAP configuration information in MashZone NextGen to ensure that email notifications are successfully delivered.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the MashZone NextGen Repositories section and click User Repository - LDAP.
3. Click Advanced Options and scroll down to find MashZone NextGen Query Properties.
4. Change the User Email address property to the name of the LDAP user attribute that contains user email addresses.
5. You may also want to update the list of attributes in Export User Attributes as MashZone NextGen Attributes and Attributes Used in Wildcard Search fields as these properties commonly include the mail attribute.
6. Click Save the changes.
7. Restart the MashZone NextGen Server to apply these updates.

8.4.11 BigMemory for Caching, Connections and In-Memory Stores

By default MashZone NextGen uses local memory for caching. This uses BigMemory as a local client that is installed with MashZone NextGen and requires only your MashZone NextGen license. In specific cases, you must also install BigMemory Servers on one or more additional hosts and

configure MashZone NextGen and the Integrated MashZone Server to work with them. MashZone NextGen requires BigMemory Servers with a BigMemory license to support:

- Significant, extensible amounts of memory, most commonly for very large datasets used with MashZone NextGen.
BigMemory Servers can be deployed in clusters, also known as Terracotta Server Arrays, that can easily be extended for scalable memory requirements.
- Distributed caching when MashZone NextGen is deployed in clusters.
With clusters, some of the internal MashZone NextGen caches must be distributed and managed by BigMemory Servers.
- Access to off-heap memory
BigMemory Servers also can manage memory outside of heap both for better scalability and performance improvements.
- Access to In-Memory Stores created and populated dynamically by external systems.
BigMemory manages the In-Memory Stores created dynamically by other systems and makes connection information available to MashZone NextGen through the Terracotta Management Console (TMC) to allow MashZone NextGen to work with this data. Apama, for example, dynamically creates distributed stores for the Apama MemoryStore which MashZone NextGen can connect to and query.

MashZone feeds that use BigMemory connections as a data source.

See [Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores](#) (page 170) for instructions on configuring MashZone NextGen to work with BigMemory Servers.

You also need to provide configuration and connection information for In-Memory Stores that are created by other systems. For more information on these tasks, see:

- [Declare BigMemory Stores for MashZone NextGen](#) (page 171)
- [Manage Dynamic BigMemory Stores for MashZone](#) (page 175).

See [Caching for the MashZone NextGen Server](#) (page 168) for an overview of MashZone NextGen caching. For caching configuration, see [Distributed Caching for MashZone NextGen Clusters](#) (page 169).

8.4.11.1 Caching for the MashZone NextGen Server

The MashZone NextGen Server caches artifact metadata, for internal purposes. By default, caches are stored in local memory. See [Artifact Caching](#) (page 169) and [Response Caching](#) (page 169) for details.

All MashZone NextGen caches can be distributed when MashZone NextGen is deployed in clusters. See [Distributed Caching for MashZone NextGen Clusters](#) (page 169) for an overview of distributed caching for MashZone NextGen.

8.4.11.1.1 Artifact Caching

Artifact caching caches metadata for feeds and other internal operational data when they are run in MashZone NextGen.

Because updates to artifacts are typically infrequent, this cache is long-lived. It is not persisted to disk. Cache entries are flushed only when an artifact is updated or when the server hosting the cache is restarted. No additional configuration is required to enable local artifact caching for MashZone NextGen.

8.4.11.1.2 Response Caching

Response caching caches the responses from dashboards and data feeds when they are run. This is a short lived cache that caches response data based on the unique signature of each request to dashboards and data feeds.

Configuration for response caching gives you fine grained control for which dashboards and data feeds use response caching and the expiration periods for their cache entries.

8.4.11.1.3 Distributed Caching for MashZone NextGen Clusters

When MashZone NextGen is deployed in clusters, artifact caching must be distributed to maintain cache integrity. Response caching, however, can be left in local memory for each MashZone NextGen Server or it can be distributed.

In many environments, local caching provides both good performance and acceptable cache integrity for response caching. Local caching is "eventually consistent", but can result in visible differences as cached responses are not guaranteed to be identical for different cluster members. For environments that cannot tolerate any loss of cache integrity, distributed response caching is recommended.

Distributed caching is only available if you purchase and deploy BigMemory Servers.

You use BigMemory Servers to handle distributed caching for MashZone NextGen.

With BigMemory Servers, data for most MashZone NextGen caches can use the total off-heap memory configured for the cluster plus any heap and off-heap memory configured for the MashZone NextGen local host.

The BigMemory Servers manage consistency and memory across the cluster. They also support failover, with mirror servers, for high availability and many other advanced capabilities that may be useful for enterprise production systems.

To configure distributed caching, see [Configure BigMemory Servers for MashZone NextGen Caching](#) (page 170) for set up instructions.

8.4.11.1.4 Configure BigMemory Servers for MashZone NextGen Caching

You can configure MashZone NextGen to work with one or an array of BigMemory Servers to provide additional memory, provide reliability and support specific other features. See BigMemory for Caching, Connections and In-Memory Stores (page 167) for more information on features that require BigMemory servers.

Procedure

1. Copy your license for BigMemory to MashZone NextGen and update MashZone NextGen startup scripts:
 - a. Copy the license file terracotta.key to the MashZoneNG-install/apache-tomcat/conf folder.
If MashZone NextGen is deployed in a cluster, you must copy this file to every cluster member.
 - b. Add the following Java system property to the MashZone NextGen server configuration file <MashZone NextGen installation>/apache-tomcat/conf/wrapper.conf:
wrapper.java.additional.<n+1>=Dcom.tc.productkey.path=MashZoneNG-install/apache-tomcat/conf/terracotta.key
Where n is the number of last additional Java parameter.
If MashZone NextGen is deployed in a cluster, you must update the server configuration files for every cluster member.
2. Edit the **ehcache.xml** file for MashZone NextGen in a text editor of your choice.
3. Find the line in ehcache.xml with <terracottaConfig> that is commented out like this:
<!-- <terracottaConfig url="localhost:9510" /> -->
Remove the comment markers and change the **url** attribute to the host (or IP address) and port for the BigMemory server(s) you installed. For example:
<terracottaConfig url="tcHost1:9510" />
There are several ways to identify one or more BigMemory servers for MashZone NextGen. See BigMemory documentation (<http://terracotta.org/documentation/>) for more information.
4. Find the line in ehcache.xml with <terracotta> that is commented out and uncomment this line for each of the following named <cache> elements:
SEARCH_RESULTS_CACHE = one part of the MashZone NextGen Artifact cache.
SERVICES_BY_ID_CACHE = one part of the MashZone NextGen Artifact cache.
SERVICE_RESPONSE_CACHE = the MashZone NextGen Response cache for dashboards and data feeds. This is optional. Update this cache **only** if you want it to be distributed.
This **<terracotta>** element allows the In-Memory Store and MashZone NextGen caches to use heap and off-heap memory in both the local host and BigMemory hosts. This combined memory is managed by BigMemory.

For more information on the `<terracotta>` element, see **Distributed Configuration** topics in BigMemory documentation (<http://terracotta.org/documentation/>).

5. Save these changes to ehcache.xml.

For clusters where this configuration file is not stored in a shared external folder, copy this file to the same location for each MashZone NextGen cluster member.

6. Start BigMemory Server(s).
7. If needed, adjust memory configuration for the local MashZone NextGen host. See Memory Configuration for the MashZone NextGen Server (page 148) for instructions.
8. Restart MashZone NextGen. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.4.11.2 Working with BigMemory Stores for RAQL

Declaring In-Memory Stores for MashZone NextGen allows MashZone NextGen to connect to stores in BigMemory that may have been created and populated by other systems. These In-Memory Stores allow you to store large datasets for quick access and to retrieve them for analysis using RAQL.

Each In-Memory Store holds one dataset with any number of rows, also known as entries, within the memory constraints defined for the store. In previous releases you could store many different datasets in one store with a dataset being one entry in the store.

To allow RAQL to work with external data, external memory stores must be configured with:

- The name of the cache manager that manages memory for the in-memory store.
Note: BigMemory does not require cache manager names, but they are a best practice for caches used as RAQL In-Memory Stores. Cache manager names prevent potential name collisions for stores.
- Search attributes for the dataset that identify the columns in the dataset and the datatype of the data in each column.

There are two types of In-Memory Stores that differ in the way they are configured in MashZone NextGen: declared and dynamic In-Memory Stores.

8.4.11.2.1 Declared In-Memory Stores

Declared stores are defined in BigMemory configuration files (ehcache.xml files) that MashZone NextGen administrators add to the MashZone NextGen Server to allow you to work with these stores.

Declared In-Memory Stores can contain data loaded by external systems. With declared in-memory stores, the producer system that stores data in the store is the system that creates the store. MashZone NextGen is always the consumer system that retrieves data from the store for analysis using RAQL.

Declaring stores before use allows you to:

- Customize store properties for each store, providing better memory management and data retention.
- Define search attributes for the columns in the dataset in each store.
- Column search attributes are required to allow RAQL to work with datasets stored by external systems.
- Defining search information also allows RAQL to delegate filtering and sorting to the In-Memory Store which provides better performance.

Configuration for declared stores specifies the cache manager, search attributes and connection information for the store. For an example, see [Declare a New In-Memory Store](#) (page 172).

8.4.11.2.1.1 Declare a new In-Memory Store

Define configuration for one or more In-Memory Stores in a cache configuration file for BigMemory (an ehcache.xml file).

It is a best practice to change the default file name **ehcache.xml** for this configuration file to something more meaningful, such as **myCRM-cache.xml**. This makes it easier to identify when multiple configuration files are uploaded to MashZone NextGen.

Procedure

1. Add a name attribute to the **<ehcache>** element and assign a unique name.

This is the cache manager name which must be unique for this MashZone NextGen Server. It should consist solely of letters, numbers, underscores(_) or dashes (-).

2. Add a **<cache>** element for each store you need to declare. The following example shows some common properties. See BigMemory documentation (<http://terracotta.org/documentation/>) for more information.

You can find this example configuration file, **sample-cache.xml**, for declared stores in the **<MashZone NextGen installation> /prestocli/raql-samples** folder.

```
<ehcache name="sample-cache" > <diskStore path="java.io.tmpdir"/> ... <cache name="StocksDeclCache" maxBytesLocalHeap="50M" memoryStoreEvictionPolicy="LRU" timeToIdleSeconds="0" timeToLiveSeconds="0"> </cache> ... </ehcache>
```

If this In-Memory Store will be populated by external systems with datasets that are Java objects, add **<searchable>** to the **<cache>** element and define a **<searchAttribute>** with the name, datatype and extractor class for each property in these Java objects that will contain data.

For the **class** attribute, use the **net.sf.ehcache.search.AttributeExtractor** interface provided in the BigMemory Search API or an implementation class of **AttributeExtractor**. See BigMemory documentation (<http://terracotta.org/documentation/>) for details.

MashZone NextGen is only able to access searchable attributes of datasets stored by external systems. For Apama used as external system, search attributes are no more required.

Since version 9.9 MashZone NextGen supports the native Apama RowValue format. MashZone NextGen can consume RowValues stored by Apama and convert them into the RAQL record format. In case of caches written by Apama searchable attributes are no more needed for accessing the data at all but they are still required for processing filters, aggregations and sorting directly in BigMemory.

3. Save this file.
4. Copy the JAR file containing the classes used as search attributes to extract data from the dataset in this cache to MashZoneNG-config/lib.

See documentation for the external system that created this dynamic store to determine what JAR files are needed. For Apama, see documentation on the MemoryStore.

The default location for this folder in MashZone NextGen is <MashZone NextGen> installation/apache-tomcat/mashzone/WEB-INF/lib. If MashZone NextGen is deployed in a cluster, however, this location may be a separate external folder. For more information, see [Setting Up an External MashZone NextGen Configuration Folder](#) (page 213).

5. Restart the MashZone NextGen Server. For instructions, see [Start and Stop the MashZone NextGen Server](#) (page 113).
6. Open the MashZone NextGen Admin Console (page 95).
7. Click **Server** to expand this section of the **Administration** menu.
8. Click **BigMemory**.
9. Open the **BigMemory Cache** tap.

The **Big Memory Cache** tab lists any In-Memory Store configuration files that have already been upload.

10. Click **Register a new EhCache Configuration File**.
11. Enter the name assigned to <ehcache> in this configuration file (in step 1) as the Big Memory Data Source Name. This name is used as a prefix for all stores defined in this configuration file to uniquely identify each store.

Data source names **must** be unique for this MashZone NextGen Server. They should contain only letters, numbers, underscores (_) or dashes (-).

If any of the declared In-Memory Stores for this connection have data populated by external systems, the data source name **must also** match the name assigned to the **<ehcache>** element in the configuration file.

12. Click **Browse** to find and select the Cache Configuration File ehcache.xml you created in step 1.
13. Click **Add this file**.

The file is uploaded to the MashZone NextGen Repository in the standard path bigmemory/caches/file-name and shown in the list by data source name. The URL shown is the relative path in MashZone NextGen to this resource.

Administrators can also manage resources files in the Admin Console. See [Manage Files for MashZone NextGen Features or Artifacts](#) (page 195) for more information.

8.4.11.2.1.2 Modify a Declared In-Memory Store

1. Update the configuration file for a declared In-Memory Store as needed.

For example, you may need to add configuration to allow an In-Memory Store to use memory in external BigMemory hosts when you add servers or deploy MashZone NextGen in staging or production environments.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the Administration menu.
3. Click **BigMemory**.
4. Open the **BigMemory Cache** tap.
5. The **Big Memory Cache** tab lists any In-Memory Store configuration files that have already been upload.
6. Select the existing BigMemory data source for this store and click **Delete**.
7. Add this data source with the updated configuration file. See Declare a New In-Memory Store (page 172) for instructions.

8.4.11.2.1.3 View Details for Declared In-Memory Stores

You can view configuration information for declared In-Memory Stores.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the Administration menu.
3. Click **BigMemory**.
4. Open the **BigMemory Cache** tap.
This lists any configuration files for declared In-Memory Store that have already been upload.
5. Click the title for a configuration file to see detailed information for the In-Memory Stores declared in that file.
6. To see the configuration file contents, click the URL for that file.

8.4.11.3 Dynamic In-Memory Stores

Dynamic stores are created programmatically on the fly when:

- An external system creates a store dynamically in BigMemory. These stores are also known as dynamic external stores because the external system is also the system the stores data in the store.

Note: Access to dynamic external stores requires BigMemory be installed as a server or server array. MashZone NextGen must also have access to the BigMemory license.

- To connect to dynamic external stores, a MashZone NextGen administrator must define connection configuration. See [Add an External Dynamic In-Memory Store Connection](#) (page 453) for more information. The Terracotta Management Console (TMC) must also be running to successfully connect to a dynamic external in-memory store.
- RAQL uses the connection configuration to retrieve configuration for the store from the Terracotta Management Console (TMC) that manages the BigMemory host for this store. This includes search attribute information that is required to allow RAQL to work with the columns in this dataset.
- One common example is the use of a dynamic external store to allow MashZone NextGen to work with datasets from distributed stores in the Apama MemoryStore.

With dynamic stores that hold external data, the external system must set a name for the cache manager and define search attributes programmatically when the external system creates the store, using the BigMemory API. See [BigMemory documentation](#) (<http://terracotta.org/documentation/>) for more information and examples.

Configuration defined in MashZone NextGen for the dynamic store allows RAQL to retrieve this configuration information.

8.4.11.3.1 Manage Dynamic BigMemory Stores for MashZone NextGen Analytics

You must define connections and identify configuration information for BigMemory stores that are created by and store data from external systems and then are used as In-Memory Stores in MashZone NextGen Analytics. For in-memory stores that are created dynamically by other systems, MashZone NextGen retrieves configuration and connection information from the Terracotta Management Console (TMC) that manages the host BigMemory Server.

You can also define connections to external in-memory stores that are not created dynamically. See [Declare BigMemory Stores for MashZone NextGen Analytics](#) (page 171) for more information.

For information on the requirements for in-memory stores that act as dynamic external stores for MashZone NextGen Analytics. For instructions on adding and managing external dynamic store configuration, see [Add an External Dynamic In-Memory Store Connection](#) (page 175) and [Delete External Dynamic In-Memory Store Connections](#) (page 176).

8.4.11.3.2 Add an External Dynamic In-Memory Store Connection

Procedure

1. Verify that the Terracotta Management Console (TMC) that manages the Big Memory Server hosting this dynamic store is running and that the store exists.

2. You should also verify that the dynamic store meets minimum requirements for MashZone NextGen.
3. Copy the JAR file containing the classes used as search attributes to extract data from the dataset in this store to <MashZone NextGen-installation/lib>.

See documentation for the external system that created this dynamic store to determine what JAR files are needed. For Apama, see documentation on the MemoryStore.

The default location for the target folder in MashZone NextGen is <MashZone NextGen-installation>/apache-tomcat/mashzone/WEB-INF/lib. If MashZone NextGen is deployed in a cluster, however, this location may be a separate external folder. For more information, see [Setting Up an External MashZone NextGen Configuration Folder](#) (page 213).

4. Restart the MashZone NextGen Server. For instructions, see [Start and Stop the MashZone NextGen Server](#) (page 113).
5. Open the MashZone NextGen Admin Console (page 95).
6. Click **Server** to expand this section of the Administration menu.
7. Click **Big Memory**.
8. Open the **Terracotta Management Server** tab.
This tab lists connections to any existing external dynamic In-Memory Stores.
9. Click **Register a new Terracotta Management Server**.
10. Enter a unique Big Memory data source name for this connection to the dynamic external cache that will act as an In-Memory Store.
11. Enter the domain and port, or IP address and port for the Terracotta Management Server. For example: localhost:9889.
12. Enter the Terracotta Management Server connection name.
Connection names cannot include periods (.), spaces or other common symbols or punctuation characters.
13. Enter the name of the Cache Manager for this cache. This name is assigned by the external system that created the cache in Big Memory.
Cache Manager names are a best practice for dynamic external stores. If the external system does not assign a cache manager name, Big Memory uses a default name which can lead to name collisions and errors.
14. If the TMC requires SSL for connections, change the Security type to SSL.
15. You can enter an user name and a password optionally.
16. Click **Add this external cache** to save this connection.

8.4.11.3.3 Delete External Dynamic In-Memory Store Connections

Procedure

1. Open the MashZone NextGen Admin Console (page 95).

2. Click **Server** to expand this section of the Administration menu.
3. Click **Big Memory**.
4. Open the **Terracotta Management Server** tab.
This tab lists connections to any existing **Terracotta Management Server**.
5. Click  **Delete** for the specific connection you want to delete.

8.4.12 Manage Terracotta DB connections

You can register, edit, delete and test Terracotta DB connections. Additionally, you can assign ACLs to an existing Terracotta DB connection.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the administration menu.
3. Click **Terracotta DB**.
4. Follow the procedure of the remaining steps:
 - Register Terracotta DB connections (Seite 177)
 - Edit Terracotta DB connections (Seite 178)
 - Test Terracotta DB connections (Seite 179)
 - Delete Terracotta DB connections (Seite 180)
 - Share Terracotta DB connections (Seite 179)
 - Register Terracotta DB connections (Seite 177)
 - Edit Terracotta DB connections (Seite 178)
 - Test Terracotta DB connections (Seite 179)
 - Delete Terracotta DB connections (Seite 180)
 - Share Terracotta DB connections (Seite 179)

You can share Terracotta DB connections with particular users and user groups so that these have access to the Terracotta DB dataset.

8.4.12.1 Register Terracotta DB connections

You can configure a connection to the Terracotta DB server. The alias is used to refer to that specific connection.

By having a connection to the Terracotta DB server configured, the dashboard developer can access all datasets available on that server. The datasets are provided in the corresponding **Terracotta DB** source operator. The alias is unique with respect to all Terracotta DB connections. Different connections can point to the same Terracotta DB server.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta DB**.
4. Click **Register new Terracotta DB connection**.
5. Enter a name for the Terracotta DB connection in the **Terracotta DB alias** field. The connection data is saved under this alias.
6. Enter the URI to the Terracotta DB. The format is terracotta://hostname:port.
7. Enable the **Materialized Feed support** option if you want to manage Materialized Feeds with this Terracotta DB connection.
 - a. Specify the **Off-heap resource** configured on the Terracotta DB server that you want to use to materialize the data feed results.
 - b. Terracotta DB provides a fast restart capability by writing the contents of a dataset in addition to the hard disk. If you want to use that feature, specify the name of the disk resource configured on the Terracotta DB server.
8. Click **Add connection**.

The **Terracotta DB** connection is created and listed by alias name.

Click  **Test connection** for the alias created to test the Terracotta DB server connection.

Parent topic: Manage Terracotta DB connections (Seite 177)

8.4.12.2 Edit Terracotta DB connections

You can edit the URI of an already existing connection to the Terracotta DB server. The alias is not editable.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta DB**.
4. Click  **Edit**.
5. Enter the URI to the Terracotta DB. The format is terracotta://hostname:port.
6. Enable the **Materialized Feed support** option if you want to manage Materialized Feeds with this Terracotta DB connection.
 - a. Specify the **Off-heap resource** configured on the Terracotta DB server that you want to use to materialize the data feed results.
 - b. Terracotta DB provides a fast restart capability by writing the contents of a dataset in addition to the hard disk. If you want to use that feature, specify the name of the disk resource configured on the Terracotta DB server.
7. Click **Update connection**.

Your changes are applied.

Click  **Test connection** to test the Terracotta DB server connection.

8.4.12.3 Test Terracotta DB connections

You can test the connection to the Terracotta DB server if the connection works correctly.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta DB**.
4. Click  **Test connection** for an alias to test the Terracotta DB server connection.

If the connection is used for Materialized Feeds, the provided values for off-heap and optional disk resource are tested as well.

The server connection is tested and a test result is displayed.

8.4.12.4 Share Terracotta DB connections

You can share Terracotta DB connections with particular users and user groups so that these have access to the Terracotta DB dataset and manage datasets with feed results.

Regardless of which users and user groups you have shared with, users with administration privilege can access all Terracotta DB aliases.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta DB**.
4. Click the  **Edit URL alias permissions** icon of the Terracotta DB connection you want to share.
5. Enter a term in the search field and click **Search**. Clicking on Search without any input values fetches all users and groups.
6. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
7. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

By default, the owner of the Terracotta DB connection is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.

8. Activate or deactivate the **Display** or **Usage** privileges of a user or user group.

A user or user group with the **Display** privilege can see the relevant source data from the datasets in the data feed or dashboard, including datasets used for Materialized Feeds.

A user or user group with the **Usage** privilege has access to the relevant Terracotta DB connection of the data source operator in the data feed editor. To access datasets with data feed results in the source operator, the user or user group needs the **Usage** privilege for the Materialized Feed alias and the **Usage** privilege for the Terracotta DB connection. To materialize data feed results in a dataset, the user or user group needs the **Usage** privilege for the materialized data feed alias, the **Usage** privilege for the associated Terracotta DB connection and the **Display** privilege for the corresponding data feed.

See also Share Materialized Feeds (page 190).

9. Click **OK**.

Your changes are applied.

8.4.12.5 Delete Terracotta DB connections

You can delete a Terracotta DB server connection.

Warning

A Terracotta DB server connection deleted can not be restored. Deleting a connection affects dashboards using datasets available over that connection as data input.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta DB**.
4. Click  **Delete**.

The Terracotta DB connection selected is deleted.

8.4.13 Manage data sources and drivers

Data sources combine the connection and driver information needed to work with both the MashZone NextGen Repository and with other databases. Data sources can use either JDBC connections or a JNDI connection pool.

See Add a data source (page 180), Edit, test or remove data sources (page 182) and Add or manage JDBC drivers (page 183) for instructions.

8.4.13.1 Add a data source

If you use connection pools to connect to databases, configure JNDI in your application server to enable access to the connection pools as needed.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).

2. Expand the **JDBC Configuration** tab.
3. If needed, add the JDBC driver for this database to MashZone NextGen. See Add or manage JDBC drivers (page 183) for instructions.
4. Click **Data Sources** to see a list of existing data sources.
5. Click **Add data source** to create a new data source.
6. Enter a data source name for a new data source.

Data source names may contain ASCII alphabetic characters and numbers **only**. Data source names may **not** contain any punctuation or symbols, such as periods (.), dashes (-) or underscores (_).

7. Select the appropriate driver for this datasource in the **JDBC Driver** drop-down menu.
8. In the **JDBC URL** field, enter either the URL for a JDBC connection or the JNDI name for a connection pool to connect to this data source. Common URL or JNDI forms include:

```
jdbc:mysql://hostname/databasename
```

For MySQL databases, it is **recommended** that you include the database name in data source URLs. If this information is omitted, testing the data source fails and may also cause errors with access to stored procedures.

```
jdbc:oracle:driver-type@hostname:port
```

```
jdbc:postgresql://hostname:port/database-name
```

```
jdbc:jtds:sqlserver://hostname:port;database-name
```

```
jdbc:sqlserver://hostname:port;databaseName=database-name
```

```
jdbc:sybase:Tds:hostname:port
```

```
java:context-path/jndi-resource-name or context-path/jndi-resource-name
```

9. Optionally, enter the username and password to use to connect to this database.
10. Click **Show connection pooling options** to display further options.
11. Optionally, set connection pooling options for this data source:

Initial Size = the initial number of connections to create when the pool for this data source starts up. This defaults to 0.

Maximum Active = the maximum number of connections that can be allocated at one time for this data source. This defaults to 8. Set this to -1 to remove all limits.

Maximum Wait = the maximum number of milliseconds that the pool will wait when no connections are available before failing. Defaults to -1 which is an indefinite wait.

Maximum Idle = the maximum number of connections that can be idle without connections being released for this data source. Defaults to 8. Set this to -1 to prevent any connections being released.

Minimum Idle = the minimum number of idle connections that can exist before new connections are added to the pool for this data source. This defaults to 0, indicating no new connections should be created.

Pool Prepared Statement = set this option to allow prepared statements for the database mashables that use this datasource to be pooled. This is disabled by default.

The usefulness and effect of pooling prepared statements depends on the type of database for this connection. See documentation for your database for more information or recommendations.

Validation Query = the SQL query that is used to validate connections in the pool for this datasource.

Validation Call Timeout = the number of milliseconds before a connection validation check is considered to have failed, causing the pool to invalidate and discard the connection. If you set this property to a number less than zero, validation calls do not expire, which is the default behavior.

Time Between Eviction Runs = the number of milliseconds between tests for idle connections for this datasource. This defaults to -1, which prevents all idle connection testing.

No of tests per run = the number of connections to test during any idle test run for this datasource. This defaults to 3.

Minimum Evictable Idle Time = the minimum number of milliseconds that a connection can be idle before being tested for eviction. This defaults to 30 minutes (1800000 milliseconds).

For more details on connection pooling properties, see Apache DBCP Documentation.

12. Click **Save changes**.

8.4.13.2 Edit, test or remove data sources

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **JDBC Configuration** tab.
3. Click **Data Sources** to see a list of existing data sources.

Initially, this lists the data source for the MashZone NextGen Repository and for the Snapshots repository.

4. To edit a data source, click the  **Edit** icon on the line for that data source and change properties.

See Add a data source (page 180) for information on specific data source properties.

5. To test the connection to a data source, click the  **Test** icon on the line for that data source.

6. To delete a data source, click the  **Delete** icon on the line for that data source.

Do **not** delete the data source for either the MashZone NextGen Repository or the Snapshots repository.

8.4.13.3 Share data sources

You can share data sources with particular users and user groups so that these have access to **JDBC data sources**.

You have administration privileges.

Regardless of the share, users with administration privilege can access all data sources.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **JDBC Configuration** to expand this section of the **Administration** menu.
3. Click **Data Sources**.
4. Click the  **Edit** data source permissions icon of the data source you want to share.
5. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
6. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
7. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

By default, the owner of the data source is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.

8. Activate or deactivate the Display or Usage privileges of a user or user group.
A user or user group with **Display** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Usage** privilege has access to the relevant alias in the data source operator
9. Click **Ok**.

Your changes are applied.

8.4.13.4 Add or manage JDBC drivers

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **JDBC Configuration** tab.
3. Click **JDBC Drivers**.
A list of existing drivers displays. This initially contains just the driver for the default MashZone NextGen Repository.
4. To remove a driver, click  **Delete** on the line for that driver.
Do **not** delete the driver for the MashZone NextGen Repository.
5. To add a new driver:
 - a. Click **Add new JDBC driver**.

- b. Enter a **Name** for a new driver.
- c. Enter the Java Class name for this driver.
- d. Click **Browse** and find the **JAR** file for this driver.
- e. Click **Add this JDBC driver**.

8.4.13.5 Migrate JDBC connections

With the MashZone NextGen version 9.10 release the persistence of JDBC drivers and connections have been changed. And only one type of JDBC connections is still available. The current version of MashZone NextGen supports the import of JDBC connections from MashZone legacy, Presto legacy and MashZone NextGen version 9.10.

8.4.13.5.1 Migrate JDBC configuration of Presto to MashZone NextGen

You can import the JDBC configuration of Presto (version 3.9 and 9.9) in MashZone NextGen.

Procedure

1. To export all existing JDBC configurations from Presto version 3.9 and 9.9 go to the `\prestocli\bin` folder of the Presto installation, open a dos command line and call:
pAdmin exportDatasource -l http://localhost:8080/mashzone/esd/api -f JDBCConnections_backup.zip -u Administrator -w manage -j
The created zip file contains all information about JDBC configurations of Presto, including the related drivers.
2. Copy the **JDBCConnections_backup.zip** file to the **prestocli\bin** folder in your MashZone NextGen installation.
3. Go to **prestocli\bin** folder in your MashZone NextGen installation, open a dos command line and call:
pAdmin importDatasource -f JDBCConnections_backup.zip -u Administrator -w manage -o

All Presto JDBC configurations will be imported into MashZone NextGen.

8.4.13.5.2 Migrate JDBC connections of Presto to MashZone NextGen

You can import the JDBC connections of Presto (version 3.9 and 9.9) in MashZone NextGen.

Procedure

This upgrade is relevant for MashZone legacy JDBC connections.

1. Copy all drivers located in the `jdbcdriers` folder of your Presto installation into the `jdbcdriers` folder of the MashZone NextGen installation and restart the MashZone NextGen MashZone NextGen Server .
2. Check if the MashZone tab exists in the Admin Console of Presto. If not, open the `presto.config` file located in `<Presto installation>\apache-tomcat\webapps\mashzone\WEB-INF\classes\` and set the `mashzone.administration.disabled=false` flag and restart the Presto server.
3. Export the required connections in the Presto Admin Console (Admin Console -> MashZone -> Database Connections tab. You have to export each connection separately. See Presto online help for details.

The exported JDBC connections are stored as `mzp` files, starting with **A_DATABASE...**, in the **importexport** folder of your Presto installation.

4. Copy all database related `mzp` files in the `importexport` folder of your Presto installation to the `dbconnections` subfolder of the `importexport` folder in your MashZone NextGen installation.
5. Start the MashZone NextGen server if not already done. Then go to the `runtool` folder (located under `Presto\mashzone\tools`), open a dos command line and call:

migrationtool -user Administrator -password manage -folder dbconnections

All JDBC connections from the **dbconnections** folder will be imported into MashZone NextGen.

In the MashZone NextGen Admin Console the JDBC connections are separated in two parts, the driver part and the data source part. You can find all JDBC related items in the **JDBC Configuration** tab of the Admin Console.

If you need to upgrade a connection using a JDBC driver that consists of multiple JAR files, you will have to create a new driver JAR which bundles all the individual files into one single file. After that, you will have to copy the newly created JAR file to the MashZone NextGen installation.

8.4.13.5.3 Migrate JDBC configuration of MashZone NextGen 9.10

You can import the JDBC configurations of MashZone NextGen version 9.10 in the current MashZone NextGen version.

Procedure

1. To export all existing JDBC configurations from MashZone NextGen 9.10 go to the `\prestocli\bin` folder of the MashZone NextGen 9.10 installation, open a dos command line and call:

pAdmin exportDatasource -l http://localhost:8080/mashzone/esd/api -f JDBCConnections_backup.zip -u Administrator -w manage -j

The created zip file contains all information about JDBC configurations of MashZone NextGen 9.10, including the related drivers.

2. Copy the **JDBCConnections_backup.zip** file to the **prestocli\bin** folder in your current MashZone NextGen installation.

3. Go to **prestocli\bin** folder in your current MashZone NextGen installation, open a dos command line and call:

```
pAdmin importDatasource -f JDBCConnections_backup.zip -u Administrator -w manage -o
```

All MashZone NextGen 9.10 JDBC configurations will be imported into the current MashZone NextGen version.

8.4.13.5.4 Migrate JDBC connections of MashZone legacy to MashZone NextGen

You can import the JDBC connections of MashZone legacy (versions 9.5 to 9.12) in MashZone NextGen.

See the **MashZone NextGen Migration Guide** for details.

8.4.14 Manage map files

8.4.14.1 Manage geoJSON files

GeoJSON is an open standard format designed for representing simple geographical features, along with their non-spatial attributes. It is based on JSON, the JavaScript Object Notation. You can upload GeoJSON files to MashZone NextGen to use in a custom map style.

All geoJSON files are stored in the following directory.

<MashZone NextGen

```
installation> \apache-tomcat\webapps\mashzone\hub\dashboard\assets\geomap\{tenant}\vectormaps
```

To support multi-tenancy, you must replace {tenant} in the path above by the corresponding tenant name. If there are no tenants, **default** is used in place of {tenant} (tenant name). If a tenant does not have a folder, the **shared** folder is used.

By default, MashZone NextGen is shipped with a few default geoJSON files in the following path.

<MashZone NextGen

```
installation> \apache-tomcat\webapps\mashzone\hub\dashboard\assets\geomap\default\vectormaps
```

The map styles are available in the **Vector map** (page 295) widget and can be selected there by file name in the **Template** drop-down menu.

The following map style templates are included.

- world-countries-by-name (default)
- world-countries-by-iso3
- world-continent-by-name
- us-states-by-name

If you add your own map style templates, we recommend that all regions (features) defined should be polygons - all non-polygon regions (features) are ignored. The **ID** attribute is the default identifier for each region (feature). The data assigned to the map widget must contain a column with values matching the **ID** attribute. Hence any polygon without an **ID** attribute is also ignored. In case there is no data assigned, the vector map could be used to select a region and publish such a selection using the region ID.

8.4.14.2 Manage tile server configuration files

TMS (Tile Map Service) is a protocol for serving maps as tiles, that is splitting the map up into a pyramid of images at multiple zoom levels. You can upload your own tile server configuration files to MashZone NextGen.

All tile server configuration JSON files are in the following directory.

<MashZone NextGen

installation> \apache-tomcat\webapps\mashzone\hub\dashboard\assets\geomap\{tenant}\tileservers

To support multi-tenancy, you must replace {tenant} in the path above by the corresponding tenant name. If there are no tenants, **default** is used in place of {tenant} (tenant name). If a tenant does not have a folder, the **shared** folder is used. By default, MashZone NextGen is shipped with a default tile server configuration. It enables you to use the Open Street Map tile servers, which has been included as a point of reference.

The maps are available in the **Map with markers** (page 286) widget and can be selected there in the **Base map** drop-down menu.

You can also add a **LeafletJS** compatible tile server configuration to host an Open Street Map data based tile server yourself. Please note, that you must be aware of all terms and licensing conditions that may apply when using a third-party tile server, including servers hosted by Open Street Map.

Example

The following configuration example points to Open Street Map tile servers hosted by Stamen.

```
{
  "label": "Stamen (watercolor)",
  "url": "http://{s}.tile.stamen.com/watercolor/{z}/{x}/{y}.jpg",
  "attribution": "Map tiles by <a href=\"http://stamen.com/\"
target=\"_new\">Stamen Design</a> Data by <a href=\"http://openstreetmap.org/\"
target=\"_new\">OpenStreetMap</a>",
  "subdomains": ["a", "b", "c", "d"]
}
```

8.4.15 Tune memory/caching for MashZone NextGen

In large installations with many users and/or many calculations, it makes sense to increase the Java heap memory, the sizes of the feed result caches and the number of calculation threads using the following techniques:

- Tune MashZone NextGen memory and cache configuration manually (page 188).
- Manual tuning gives you control to balance memory requirements for MashZone NextGen, but does require manual updates to several configuration files.

8.4.15.1 Tune MashZone Memory and Cache Configuration Manually

To manually update memory and cache configuration

- Update Cache Memory Settings (page 188)
- Update MashZone ThreadSize Properties (page 188).
- Then restart the MashZone NextGen Server to apply this change. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

8.4.15.2 Update Cache Memory Settings

- In the text editor of your choice, open the **ehcache.xml** file in the `web-apps-home/mashzone/WEB-INF/classes` folder.
- Update the **maxBytesLocalHeap** value on the **<cache>** elements with the following names:
 - **RESULT_FEED_BASE**
 - **RESULT_FEED_TOP**
 - **RESULT_FEED_DEBUG**
- Save your changes

8.4.15.3 Update MashZone ThreadSize Properties

- In the text editor of your choice, open the **mashzone.properties** file in the `web-apps-home/mashzone/WEB-INF` folder.
- Update the following properties:
 - **calculation.threadpool.coresize**
 - **calculation.threadpool.maxsize**
- Save your changes

8.4.16 Manage Materialized Feeds

8.4.16.1 Configure Materialized Feeds

Using Materialized Feeds (page 86) you can materialize the results of a data feed evaluation in Terracotta DB.

In order to work with those results, you can run RAQL queries over the corresponding dataset using the Terracotta DB (page 311) data source operator.

Prerequisite

You have configured a Terracotta DB connection to use materialized data feeds. (Seite 177)

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Materialized Feeds > Materialized Feeds** to expand this section of the **Administration** menu.
3. Click **Create**.
4. Enter a unique name for this Materialized Feed in the **Alias** input box. The configuration is saved under this alias.
5. Select a data feed in the **Feed** drop-down menu. The menu contains all data feeds you can access.
6. Select a Terracotta DB connection in the **Terracotta DB connection** drop-down menu. The menu contains all Terracotta DB connections supporting Materialized Feeds.
The dataset for materializing feed results will be created on the Terracotta DB server defined for this Terracotta DB connection.
7. Enter a name of the dataset used for materializing feed results in the **Dataset name** input box.
Either use the predefined name or set your own name. The name mandatory begins with the prefix **MZNG**. The dataset is created if it not already exists. Note that the dataset contents must have the same schema. Trying to insert data with a different schema into an existing dataset fails with an exception.
8. Click **Save**.

The configuration for Materialized Feeds is created and listed by alias name.

8.4.16.2 Edit Materialized Feeds

You can edit existing configurations of Materialized Feeds.

The alias and the associated data feed cannot be edited.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Materialized Feeds > Materialized Feeds** to expand this section of the **Administration** menu.

3. Click the  **Edit** icon of the configuration you want to edit.
4. Select a Terracotta DB connection in the **Terracotta DB connection** drop-down menu. The menu contains all Terracotta DB connections supporting Materialized Feeds.
The dataset for materializing feed results will be created on the Terracotta DB server defined for this Terracotta DB connection.
5. Enter a name of the dataset used for materializing feed results in the **Dataset name** input box.
A new, empty dataset will be created with the new dataset name. The old dataset will not be deleted.
6. Click **Update**.

Your changes will be applied.

8.4.16.3 Delete Materialized Feeds

You can delete existing configurations of Materialized Feeds.

Warning

A deleted configuration cannot be restored. Deleting a configuration affects dashboards using the dataset with data feed results as data input.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Materialized Feeds > Materialized Feeds** to expand this section of the **Administration** menu.
3. Click the  **Delete** icon of the configuration you want to delete.
4. In the confirmation dialog you can decide if you want to delete the associated dataset. Enable the corresponding option.
5. Click **Yes**.

Your selected configuration is deleted.

8.4.16.4 Share Materialized Feeds

You can share Materialized Feeds with particular users and user groups so that these can work with Materialized Feed results.

You must grant the appropriate privileges to the users or user groups with whom you want to share the Materialized Feeds.

Regardless of which users and user groups you have shared with, users with administration privilege can manage all Materialized Feeds.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Materialized Feeds > Materialized Feeds** to expand this section of the **Administration** menu.
3. Click the  **Edit alias permission** icon for the Materialized Feed you want to share.
4. Enter a term in the **Search** input box and click the  **Search** icon. Clicking the **Search icon** without any input values fetches all users and groups.
5. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
6. Drag an user or user group from the **Search results** field and drop it into the **Principals with permissions field**.
By default, the owner of the Materialized Feed is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.
7. Enable or disable the **Display** or **Usage** privileges of a user or user group.
8. Click **OK**.

Your changes are applied.

In addition to configuring the sharing settings for Materialized Feeds, you must also configure the required privileges for the data feed and the Terracotta DB connection used. For details, see [Which privileges are required for Materialized Feeds? \(page 89\)](#)

8.4.16.5 Materialize data feed results in Terracotta DB

When you have configured a Materialized Feed, you can trigger the calculation of data feed results and the subsequent materialization of the results in the corresponding Terracotta DB dataset.

That process cannot be run multiple times in parallel for the same Materialized Feed. If you want to trigger a new run, you must wait until the current run is completed. Each time a new run is started, the contents of the dataset are deleted so that it contains only the results of the latest run.

You have various options for materializing data feed results.

8.4.16.5.1 Use the Admin Console

You can materialize data feed results in Terracotta DB using the MashZone NextGen Admin Console.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Materialized Feeds > Materialized Feeds** to expand this section of the **Administration** menu.

3. Click the  **Upload** icon to calculate and materialize the data feed results.

For feeds with input parameters, the default values of those input parameters are used for feed execution.

4. Click **Yes** to confirm.

The data feed results are calculated and materialized using the configured Terracotta DB connection.

After the feed calculation and the subsequent materialization of results are completed, you can access the feed results using the Terracotta DB (page 311) source operator.

Tip

Additional information like duration or potential execution errors is available in the MashZone NextGen log file.

8.4.16.5.2 Use the API Console

You can materialize data feed results in Terracotta DB using the MashZone NextGen API Console.

Procedure

1. Open the API Console. (page 438)
2. Under **Pick an API** click **MaterializedFeedService** > **materializeFeed**.

In the request template adapt the alias. For the case of feeds with input parameters, specify the list of input parameters of the data feed and their values to be used for calculating the data feed results. The list is represented in JSON as name/value pairs, for example, `{"param1\": \"ABC\", \"numberList\": [\"123\", \"456\"]}`. If the input parameters are not defined, default values as defined in the data feed are used.

3. Click **Run**.

The data feed results are calculated and materialized using the configured Terracotta DB connection.

After the feed calculation and the subsequent materialization of results is completed, you can access the feed results using the Terracotta DB (page 311) source operator.

Tip

Additional information like duration or potential execution errors are available in the MashZone NextGen log file.

8.4.16.5.3 Use the server administration tool

You can materialize data feed results in Terracotta DB using the MashZone NextGen server administration tool.

You can use the server administration tool for scheduling the materializing procedure.

Procedure

1. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** directory.

2. Run the following command.

```
padmin materializeFeed -a alias -i input-parameters [-l prestoURL] -u username -w password
```

-a alias: Is the alias of the Materialized Feed configuration.

-i input-parameters: Is the list of input parameters of the data feed and their values to be used for calculating the data feed results. The list is represented in JSON as name/value pairs, for example, {"param1": "ABC", "numberList": ["123", "456"]}. If the input parameters are not defined, default values as defined in the data feed are used.

-u username: Is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.

-w password: Is the MashZone NextGen password to log in with.

The data feed results are calculated and materialized using the configured Terracotta DB connection.

General messages and errors from the export process are sent to the command window (stdout).

After the feed calculation and the subsequent materialization of results is completed, you can access the feed results using the Terracotta DB (page 311) source operator.

Tip

Additional information like duration or potential execution errors are available in the MashZone NextGen log file.

8.4.16.5.4 Use server web services

You can materialize data feed results in Terracotta DB using the MashZone NextGen server web services. You can use a public REST web service such as cUrl. In addition, you can combine those calls with an external scheduling tool, such as cron jobs, to schedule the executions.

Procedure

1. You can use the public login service to authenticate and set a cookie.

Authentication using basic credentials, see below.

2. Using that authentication cookie, you can trigger the data feed calculation and materialization using

- Relative service URL: /mashzone/mzservices/tcstore/materializefeed

- Method: POST

- Parameter: name=materializedFeed, value="the name of the Materialized Feed alias to use"
- Parameter: name=inputParameters, value="[{"param1": "someValue", "param2": "anotherValue"}]"

3. Using above authentication cookie, you can use the public logout service to finish your session.

Authentication using basic credentials, see below.

The data feed results are calculated and materialized using the configured Terracotta DB connection.

After the feed calculation and the subsequent materialization of results is completed, you can access the feed results using the Terracotta DB (page 311) source operator.

Tip

Additional information like duration or potential execution errors are available in the MashZone NextGen log file.

AUTHENTICATION USING BASIC CREDENTIALS

For authentication using basic credentials, you must use the REST API method to login to create a MashZone NextGen session and authenticate the user. This adds a cookie with session information which is then included in all subsequent requests. You can also use the REST API method to logout and end the session or simply let the session expire (based on MashZone NextGen timeout configuration).

- **Login:** use this URL:

```
http://app-server:port/mashzone/edge/api/rest/UserManagerService/login?presto_username=username&presto_password=pw
```

Both parameters must be in clear text.

Successful logins set a session cookie and return an XML response in the form:

```
<?xml version="1.0"?>
<ServiceResponseObject>
  <version>2.0</version>
  <appId></appId>
  <sid>UserManagerService</sid>
  <svcVersion>1.0</svcVersion>
  <oid>login</oid>
  <response>
    <authToken>Token is set as HTTP cookie</authToken>
  </response>
  <errorCode></errorCode>
  <invId></invId>
  <header/>
  <serviceHeaders/>
</ServiceResponseObject>
```

Unsuccessful logins return a simple text response: **Authentication failed.Userid:someName.**

- **Logout:** use this URL:

```
http://app-server:port/mashzone/edge/api/rest/UserManagerService/logout
```

8.5 MashZone NextGen Server Administration

8.5.1 Manage Files for MashZone NextGen Features or Artifacts

MashZone NextGen uploads and hosts files for dashboards and data feeds that are not accessible via HTTP (spreadsheets, CSV or XML files). These files are saved and managed in the MashZone NextGen Repository to ensure better management of resources and easier deployment or migration across different environments and versions.

MashZone NextGen administrators may need to manually add files to MashZone NextGen to provide data files or resources for dashboards or data feeds.

Common management tasks for files include: Add External Resources as MashZone NextGen Files (page 195), Find MashZone NextGen Files (page 196) or Update or Delete MashZone NextGen Files (page 196).

8.5.1.1 Add External Resources as MashZone NextGen Files

You can add external resources to MashZone NextGen to make them easily accessible in dashboards and data feeds.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the Platform Features section and click **File Resources**.
3. Click **Upload New Files**.
4. Click **Browse**, find and select the file you want to upload and click **Open**.

The location and file name fill in and a new set of fields open to upload another file.

5. If needed, add to the path or change the file name.

The name of the file defaults to **/file-name**. **If you accept the default, the URL to access this file becomes `http://app-server:port/mashzone/files/file-name`.**

You can organize files into 'pseudo folders' by adding to the path, using **/** as the separator. For example, a file name of **/images/reports.png** has a URL of

`http://app-server:port/mashzone/files/images/reports.png` and can be found in file search (along with any other files in the 'images folder') by searching for **images** as the file name.

You can also upload files that are normally loaded automatically, such as thumbnails. Simply specify the standard path.

6. Repeat the steps, as needed, to find and name any other files you want to upload.

7. Click  **Delete** to close the empty file upload fields.
8. Click **Upload files**.

The files are added to the MashZone NextGen Repository and are now available via a MashZone NextGen URL.

8.5.1.2 Find MashZone NextGen Files

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Platform Features** section and click **File Resources**.
3. Enter either:
 - Part of the file name(s).
 - Part of the path to the file(s).
4. Click **Search**.
 - File search results are always sorted by path and file name.

8.5.1.3 Update or Delete MashZone NextGen Files

Although rare, you may occasionally need to update or even delete files from MashZone NextGen.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Platform Features** section and click **File Resources**.
3. Find the specific file you need to update or delete. See Find MashZone NextGen Files (page 196) for techniques.
4. To upload an updated file:
 - a. Click **Edit** on the line for that file.
 - b. Click **Browse** and find the updated file you want to replace the existing file in MashZone NextGen.
 - c. Click **Upload this file**.
5. To delete a file, click **Delete** on the line for that file.

8.5.2 Manage resource directories

Resource directories hold file-based data sources, such as Excel spreadsheets, CSV or XML files. The resource aliases can be used by the data source operators to read local files.

8.5.2.1 Create resource directory

To work with data sources in MashZone NextGen Feed Editor that are file-based, such as Excel spreadsheets, CSV files or XML files, you must store the files in a **resource directory** that the Integrated MashZone NextGen Server knows. This can be the default resource directory:

MashZoneNG-install/mashzone/data/resources

Or it can be any subdirectory of the default.

You can also use resource directories to control access to data source files to specific users or groups.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **File resource** -> **File resource** to open the resource alias page.
3. Click **Create**.
4. Give the directory an alias name of your choice in the **Resource directory** input box.
You cannot modify the alias name later.
5. Enter the Path of the new resource directory.
6. Click **Add resource**.

The new resource directory is created and is displayed in the list with the specified alias.

8.5.2.2 Change resource directory

You can adapt the path of already existing resource directories.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **File resource** -> **File resource** to open the resource alias page.
3. Click the  **Edit resource alias** icon of the resource you want to edit.
4. Enter the **Path** of the resource directory.
5. Click **Save resources**.

Your changes are applied.

8.5.2.3 Delete resource directory

You can delete existing resource directories.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click File resource -> File resource to open the resource alias page.
3. Click the  **Delete resource alias** icon of the resource you want to delete.

4. Click **Yes**.

The directory selected is deleted from the list.

8.5.2.4 Share resource directory

You can share resource directories with particular users and user groups so that these have access to the directory content.

You have administration privileges.

Regardless of the share, users with administration privilege can access all resource directories.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Resource alias** -> **Resource alias** in the **Administration** menu.
3. Click the  **Edit resource permissions** icon of the resource you want to share.
4. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
5. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
6. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.
By default, the owner of the resource directory is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.
7. Drag an user or user group from the Search result field and drop it into the Principals with permissions field.
By default, the owner of the **Apama Event Services** is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.
8. Click **Ok**.

Your changes are applied.

8.5.3 Manage URL aliases

You can manage your URL aliases in the **Admin console**.

Using an URL alias is always recommended to shorten the link used in ,for example, dashboards and data feeds. You have to enter the path where the data are stored only, and not the complete URL. The resource aliases can be used by the data source operators to read local files.

8.5.3.1 Create URL alias

You can create URL aliases to shorten a link used in ,for example, dashboards and data feeds.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **URL aliases** -> **URL aliases** to open the **URL alias** page.
3. Click **Create**.
4. Give the URL an alias name of your choice in the **Alias** input box.
You cannot modify the alias name later.
5. Enter the URL.
6. Activate the **Use basic authentication** option if an authentication is require for using the URL.
 - a. Enter a user name.
 - b. Enter the password associated with the user name.
7. Click **Add alias**.

The new URL alias is created and is displayed in the **URL alias** list.

8.5.3.2 Change URL alias

You can adapt the URL and the authentication credentials of already existing URL aliases.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **URL aliases** -> **URL aliases** to open the **URL alias** page.
3. Click the  **Edit URL alias** icon of the URL alias you want to edit.
4. Make your changes.
5. Click **Save alias**.

Your changes are applied.

8.5.3.3 Delete URL alias

You can delete existing URL aliases.

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **URL aliases** -> **URL aliases** to open the **URL alias** page.
3. Click the  **Delete URL alias** icon of the URL alias you want to delete.
4. Click **Yes**.

The URL alias selected is deleted from the list.

8.5.3.4 Share URL alias

You can share URL aliases with particular users and user groups so that these have access to the directory content.

You have administration privileges.

Regardless of the share, users with administration privilege can access all URL aliases.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **URL aliases** -> **URL aliases** to open the **URL alias** page.
3. Click the  **Edit URL alias permissions** icon of the alias you want to share.
4. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
5. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the Search results field.
6. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

Note: By default, the owner of the URL alias is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.

7. Activate or deactivate the **Display** or **Usage** privileges of a user or user group.
A user or user group with **Display** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Usage** privilege has access to the relevant alias in the data source operator.
8. Click **OK**.

Your changes are applied.

8.5.4 Deploying MashZone NextGen Instances, Clusters or Artifacts

Deploying MashZone NextGen to new hosts or new environments typically involves Deploying the Core Widgets (page 200), shown below, and optionally Deploying MashZone NextGen Artifacts and Other Metadata (Seite 201).

8.5.4.1 Deploying the Core Widgets

The core widgets include the MashZone NextGen Server, the MashZone NextGen Repository, which is typically installed in a database other than the default Derby database, and the MashZone NextGen Analytics In-Memory Stores and MashZone NextGen caches.

In earlier releases, the MashZone NextGen Hub and AppDepot were deployed in a separate web application from the MashZone NextGen web application. Effective in 3.2, all the core widgets are deployed in the single web-apps-home/presto web application.

For individual MashZone NextGen servers, you typically do a default installation (see [Installing Software AG Products](#)). You may also move the MashZone NextGen Repository to a database of your choice. See [Move the MashZone NextGen repository to a robust database solution](#) (page 116) for instructions.

You can leave the the MashZone NextGen Analytics In-Memory Stores and MashZone NextGen caches in local memory for a single MashZone NextGen server. This uses the default client installation of BigMemory. If additional memory or reliability is required, you can also deploy BigMemory as an add-on in a separate host or cluster. See [Working with MashZone NextGen Analytics In-Memory Stores](#) for more information and links.

To deploy multiple unclustered servers, see [Deploying Multiple MashZone NextGen Servers in One Host](#) (page 209). To deploy MashZone NextGen servers in clusters, see [Clustering MashZone NextGen Servers](#) (page 209) for requirements and links.

8.5.4.2 Deploying MashZone NextGen Artifacts and Other Metadata

You deploy specific artifacts and metadata from a source MashZone NextGen Server to a target MashZone NextGen Server using the export and import commands.

Important: You **cannot** use export and import commands when the MashZone NextGen version for the source and target MashZone NextGen Servers are different:

- For major upgrades, use the migrate command instead.
- For minor upgrades, please contact Technical Support or your Software AG representative.

In addition to the basic metadata for an artifact, a successful deployment must include related metadata, related files, extensions the artifact may use and any other artifacts that the artifact depends on.

The export and import commands automate deployment for most of this data, with some specific limitations that require manual deployment steps.

Procedure

1. Export the specific artifacts that you want to deploy to another MashZone NextGen Server and any artifacts that they may use.

See the following topic for instructions using these MashZone NextGen export commands:

[Export users, user metadata and groups](#) (page 204)

[Export dashboards](#) (Seite 77)

[Export data feeds](#) (Seite 205)

2. Copy the files for any extensions used by the exported artifacts from the MashZoneNG-config folder for the source MashZone NextGen Server to the MashZoneNG-config folder for the target MashZone NextGen Server.

Note: The MashZoneNG-config folder may be an external configuration folder outside of the source and target MashZone NextGen Servers or it may be in the default locations. See [Setting Up an External MashZone NextGen Configuration Folder \(page 213\)](#) for more information on MashZoneNG-config locations.

3. Define datasources in the Admin Console for the target MashZone NextGen Server with matching names and JDBC drivers to the datasources in the source MashZone NextGenMashZone NextGen Server.

See [Manage data sources and drivers \(page 180\)](#) for instructions.

4. Use the export files created earlier to import dashboards and data feeds, users, groups and user group assignments from the source MashZone NextGen Server.

See the following topics for information on using these commands:

[Import dashboards \(Seite 78\)](#)

[Import data feeds \(Seite 208\)](#)

[Import users, user metadata and groups \(page 206\)](#)

Table 1. Known Export/Import Limitations

	Exported	Not Exported
Artifact Metadata	<ul style="list-style-type: none"> ▪ Basic metadata such as provider, category, tags and description. ▪ Ownership (who created the artifact). ▪ On/off status. ▪ Run permissions. ▪ Views. ▪ Artifact attributes. ▪ For apps, the AppDepot status. 	<ul style="list-style-type: none"> ▪ User rating and feedback. ▪ Snapshots. ▪ Snapshot schedules. ▪ Caching configuration.
Related Metatdata/ User Metadata	<ul style="list-style-type: none"> ▪ Providers. ▪ Categories. ▪ Global and user MashZone NextGen attributes. ▪ Users, groups and user group assignments if this data is tracked in the default MashZone NextGen User Repository and not in your LDAP Directory. 	<ul style="list-style-type: none"> ▪ Datasources and their JDBC drivers that are used by database mashables or by mashups. ▪ Datasources must be added to the target MashZone NextGen Server before you import any artifacts that use them or the import will fail. ▪ For apps that are published to the AppDepot, any user preferences for Favorite Apps. ▪ User preferences for apps in

		Mashboard or the Mashboard state for workspace apps.
Resource Files	<ul style="list-style-type: none"> ▪ Thumbnails for apps or pluggable views. ▪ Screenshots for apps. ▪ HTML, JavaScript, CSS, images or any other file uploaded in the package for a custom app. 	Thumbnails for mashables or mashups.
Dependencies	<p>Optionally can export dependent artifacts:</p> <ul style="list-style-type: none"> ▪ For workspace apps, this always exports all the apps used in the workspace. ▪ For individual apps, you can choose to also export any mashables or mashups explicitly declared and used by those apps. ▪ If you choose to include dependencies, all dependencies for basic apps are handled because MashZone NextGen automatically declares dependencies for basic apps. ▪ For custom apps, export handles any dependencies that are explicitly declared with a <dependson> element in the App Specification. App developers must supply this information. ▪ For mashups, this always exports any other mashups or mashables that are used by the mashup. ▪ For pluggable views or pluggable libraries you can choose to export any library dependencies. 	Any snapshots used by apps.
Extensions	<ul style="list-style-type: none"> ▪ Registered macros, for use in mashups. 	<ul style="list-style-type: none"> ▪ Attribute definitions for extension attributes in

	<ul style="list-style-type: none"> HTML, JavaScript, CSS, images or any other file uploaded in the package for a pluggable view or pluggable library. 	<ul style="list-style-type: none"> artifacts. Any of the file-based extensions such as custom XPath functions. See Setting Up an External MashZone NextGen Configuration Folder (page 213) for a complete list.
MashZone NextGen Server Configuration		Configuration for the MashZone NextGen Server.

8.5.4.3 Export users, user metadata and groups

You can export all users, user groups, user group assignments from the MashZone NextGen Repository to an export file. You can then import this file to another MashZone NextGen Repository.

Important: If you have configured MashZone NextGen to work with your LDAP Directory, this command **only** exports MashZone NextGen User Attributes. Data for users, user groups and user group assignments resides in LDAP.

Procedure

- If it is not running, start the MashZone NextGen Server for the MashZone NextGen Repository with the user groups that you wish to export. See Start and Stop the MashZone NextGen Server (page 113) for instructions.
- Open a command window and move to the MashZoneNG-install/prestocli/bin folder.
- Enter this command:

```
padmin exportUsersRoles -f output-file [-l prestoURL] -u username -w password [-v]
```

-f output-file: is the path and name of the export file to hold the metadata.

-l prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to **http://localhost:8080/mashzone/edge/api**.

-u username: is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.

-w password: is the MashZone NextGen password to log in with.

-v: is an optional flag to turn on verbose logging.

All messages and errors from the export process are sent to the command window (stdout). Once the export command completes successfully, you can use the output file to import the data to another MashZone NextGen Repository.

8.5.4.4 Export dashboards

You can export your MashZone NextGen dashboards.

You can use the export zip file to create a backup or to import your dashboards into another MashZone NextGen installation.

Procedure

1. Open a command window and move to the **MashZoneNG-install/prestocli/bin** folder.
2. Enter this command:

```
padmin exportDashboard -i identifier [-f output-file] [-l prestoURL] -u username -w password [-v] [-o]
```

- -i identifier: Mandatory dashboard identifier. It can be "id=", "name=" or "all", enclosed in quotes.
- -i "name=dashboardname": If there are multiple dashboards with the same name only the first dashboard found will be exported.
- -i "id=43243244434432": The dashboard ID (GUID) is unique in the MashZone NextGen system.
- -i "all": Exports all dashboards for that user.
- -f output-file: Optional path and name for the export. If omitted, an output zip file is created in the folder in which this command is executed:
Single export with option -i "id=3456" or "name=name" create a new file with name "name_guid.zip"
Multiple export with option -i "all" create a new file dashboard-export-timestamp.zip
- -l prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to **http://localhost:8080/mashzone/esd/api**.
- -u username: **MashZone NextGen user name to log in with. This account must have MashZone NextGen administrator permissions.**
- -w password: is the MashZone NextGen password to log in with.
- -v: is an optional flag to activate verbose logging.
- -o: Optional flag to overwrite an existing export file.

Once the export command completes successfully, you can use the output file to import dashboards into MashZone NextGen.

Permissions for each dashboard were automatically stored in the zip file. If no permissions are assigned to the dashboard, the permission file saved is empty.

The zip file also includes information about the dashboard creator.

8.5.4.5 Export data feeds

You can export your MashZone NextGen data feeds.

Export creates an export file that you can use to import data feeds to MashZone NextGen.

Procedure

1. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.

2. Enter this command:

```
padmin exportFeed -i identifier [-f output-file] [-l prestoURL] -u username -w password [-v] [-o]
```

-i identifier: mandatory data feed identifier. It can be "id=", "name=" or "all", enclosed in quotes.

-i "name=feedname": If there are multiple data feeds with the same name then only the first founded data feed will be exported.

-i "id=43243244434432": The data feed id (Guid) is unique in the MashZone NextGen system.

-i "all": Export of all data feeds for that user.

-f output-file: an optional path and name for the export file to put data feeds. If omitted, this generates an output zip file in the folder where this command is executed:

Single export with option -i "id=3456" or "name=name" create a new file with name "name_guid.zip"

Multiple export with option -i "all" create a new file datafeed-export-timestamp.zip

This file must not already exist, unless you also use the **-o** option.

-l prestoURL: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to **http://localhost:8080/mashzone/esd/api**.

-u username: is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.

-w password: is the MashZone NextGen password to log in with.

-v: is an optional flag to turn on verbose logging.

-o: an optional flag to overwrite an existing export file. If you omit this option, the output file must not already exist.

Once the export command completes successfully, you can use the output file to import data feeds to MashZone NextGen.

Permissions for each data feed were automatically stored into the zip file. If there are not any permissions assigned to the data feed an empty permission file is stored.

There is also an information about the data feed creator stored in the zip file.

8.5.4.6 Import users, user metadata and groups

You must have a users export file to import. See Export users, user metadata and groups (page 204) for instructions.

This command is available only in MashZone NextGen 3.2 or later.

A user export file contains MashZone NextGen User Attributes. It may also contain users, user groups and user group assignments if you are using the default MashZone NextGen User Repository rather than an LDAP Directory.

Procedure

1. If it is not started, start the MashZone NextGen Server for the MashZone NextGen Repository where you wish to import data. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

2. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.

3. Enter this command:

```
padmin importUsersRoles -f input-file [-I prestoURL] -u username -w password [-c] [-o] [-v]
```

-f input-file: is the path and name of the export file to import data from.

-I prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to

http://localhost:8080/mashzone/edge/api.

-u username: is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.

-w password: is the MashZone NextGen password to log in with.

-c: is an optional flag to allow the import process to continue if errors occur during the import. By default, any import errors stop all further processing.

-o: is an optional flag to allow import information for a MashZone NextGen global attribute to overwrite an existing user, group, user group assignments or MashZone NextGen User Attribute with the same ID.

-v: is an optional flag to turn on verbose logging.

Messages and errors from the import process are sent to the console window (stdout). Once the import is successfully finished, you may confirm that the appropriate data has been imported in MashZone NextGen.

8.5.4.7 Import dashboards

You can import dashboards in MashZone NextGen.

The dashboards are saved in a zip containing the dashboard definition, resource policy, and dashboard permissions, etc.. If you import a dashboard including permissions, the creator of the dashboard can view and edit the dashboard. The importer of a dashboard automatically becomes the creator of the dashboard if the dashboard is imported without permissions.

Procedure

1. Open a command window and move to the **<MashZoneNG-installation>/prestocli/bin** folder.

2. Enter this command:

```
padmin importDashboard [-l prestoURL] -f input-file -p importPermissions -u username -w password [-v] [-o]
```

- -f input-file: Path and name of the import zip file.
- -p importPermissions: Imports the resource policy and permissions saved in the import zip file.
- The importer of a dashboard automatically becomes the creator of the dashboard if the dashboard is imported without permissions. And only administrators can see and work with the dashboards imported.
- -o: is optional. Allows overwriting an existing dashboard in MashZone NextGen Dashboard.
- -l prestoUrl: Optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this results in **http://localhost:8080/mashzone/esd/api**.
- -u username: is the MashZone NextGen user name to log in with. This account **must** have MashZone NextGen administrator permissions.
- -w password: is the MashZone NextGen password to log in with.
- -v: is an optional flag to turn on verbose logging.

Once the import command completes successfully, you can use the imported dashboards in MashZone NextGen.

If you have imported dashboards from Presto 3.9 into MashZone NextGen, save the imported dashboards in edit mode of the widget before you display them in view mode. Otherwise, an error message is displayed.

8.5.4.8 Import data feeds

You can import data feeds to MashZone NextGen.

The data feeds are saved in a zip file that contains among other things the data feed definition, resource policy and data feed permissions. If you import a data feed including the permissions then the creator of the data feed can view and edit the data feed. Importing data feeds without the relevant permissions makes the importer automatically to the creator of these data feeds.

Procedure

1. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.

2. Enter this command:

```
padmin importFeed [-l prestoURL] -f input-file -p importPermissions -u username -w password [-v] [-o]
```

f input-file: path and name for the import zip file.

p importPermissions Imports the resource policy and permissions saved in the import zip file.

If you import data feeds without permissions makes the importer automatically to the creator of these data feeds and the data feeds has no explicit permissions which means that only administrators can see and work with the data feeds .

o: is optional. Allows to overwrite an existing data feeds.

I prestoUrl: is optional. Use this if the MashZone NextGen Server is remote or if it is not running in Tomcat on the default Tomcat port. If you omit this option, this defaults to `http://localhost:8080/mashzone/esd/api`.

u username: is the MashZone NextGen username to log in with. This account must have MashZone NextGen administrator permissions.

w password: is the MashZone NextGen password to log in with.

v: is an optional flag to turn on verbose logging.

Once the import command completes successfully, you can use the imported data feeds in MashZone NextGen feed editor.

8.5.4.9 Deploying multiple MashZone NextGen servers in one host

You can deploy several different, independent MashZone NextGen Servers on a single host. Each MashZone NextGen Server must be hosted in its own application server and have its own MashZone NextGen Repository.

To host multiple, independent servers, simply install each being sure to change the ports assigned to each MashZone NextGen Server, MashZone NextGen Repository and the administration port for Tomcat.

Note: You can also create clusters of MashZone NextGen Servers to provide load balancing. See Clustering MashZone NextGen Servers (page 209) for information.

8.5.5 Clustering MashZone NextGen Servers

In production environments, it is common to use clustering solutions to provide better performance for various loads, to provide high availability or to provide both. Because MashZone NextGen is a web application, using an HTTP session based on J2EE standards, you can apply the same cluster architectures and solutions to MashZone NextGen that you use with other web applications.

See Setting Up a New Cluster (page 209) or Adding New Members to an Existing Clusters (page 211) for the tasks you need to complete.

8.5.5.1 Setting Up a New Cluster

The configuration and deployment of a new cluster requires these basic steps:

- **Setting Up an External MashZone NextGen Configuration Folder (page 213):** this allows you to keep most of the configuration and extensions for MashZone NextGen in a single set of folders that can be shared across the entire cluster. This simplifies both the initial configuration as well as ongoing updates and deployment of new mashables, mashups or apps.

Note: This step is highly recommended, but not required. If you do not use a shared configuration folder, all subsequent updates to configuration or extensions for new artifacts must be manually copied to each member of the cluster.

- This folder should reside in a file system that is shared or mounted across the cluster. You may also need to provide data redundancy or failover capabilities for this shared file system.
- As part of this step, you also typically deploy one MashZone NextGen Server in the cluster and complete most of the basic configuration that will be shared across the cluster.
- **Sharing the MashZone NextGen Repository in Clustered Environments (page 212):** all nodes in the cluster work with a shared MashZone NextGen Repository which you must create and configure.
- Sharing the MashZone NextGen Repository does not, by itself, provide any data redundancy, load balancing or failover capabilities for the database. These requirements are handled in the data layer by your database server or other replication/synchronization solutions, such as DRBD. For more information, see documentation for your database or replication/synchronization solution.
- **Configuring Caching for the Cluster:** each MashZone NextGen Server has a local cache for mashable and mashup responses as well as local caches for updates to artifacts. If MashZone NextGen Analytics is enabled in your MashZone NextGen license, the MashZone NextGen Analytics In-Memory Stores are also local.
- In clusters you:
 - Can leave the response cache as a local cache or you can configure a distributed cache that all MashZone NextGen Servers in the cluster share.
 - **Must** configure a distributed cache for artifact updates that all MashZone NextGen Servers in the cluster share.
 - **Must** configure a distributed cache for the MashZone NextGen Analytics In-Memory Stores that all MashZone NextGen Servers in the cluster share.
- See **Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores (page 170)** for instructions on how to configure BigMemory, or other caching solutions, as a distributed cache for .MashZone NextGen
- **Defining the Application Server Cluster:** the application servers that host each MashZone NextGen Server define and handle clustering requirements at the application layer. You can also add a load balancer to the cluster.
- In addition to the basic cluster configuration required by your application server and load balancer, MashZone NextGen has a single requirement for application-layer cluster configuration. You must either:
 - Enable session replication in each application server in the cluster.

- Enable session affinity, sometimes also called 'sticky sessions,' in the load balancer.
- Or do both.
- See documentation for your application server and/or load balancer for information on how to do this.
- **Adding Additional MashZone NextGen Servers to the Cluster:** once you have set up the shared resources, you can deploy and add additional members to the cluster. See Adding New Members to an Existing Cluster (page 211) for instructions.
- **Add MetaData and Deploy Artifacts:** for this new environment. For artifacts, you can automate some parts of this process using export and import commands. See Deploying MashZone NextGen Artifacts and Other Metadata (Seite 201) for instructions.

8.5.5.2 Adding New Members to an Existing Cluster

To add additional MashZone NextGen Servers to an existing cluster

1. Install the MashZone NextGen Server. See Installing Software AG Products for instructions.
2. Configure the MashZone NextGen Server to use the shared MashZone NextGen Repository for the cluster. See Share an Existing MashZone NextGen Repository (page 213) for instructions.
3. If the cluster has a shared external configuration folder, add this folder and any subfolders to the classpath for the MashZone NextGen Server's application server to enable access to this shared configuration.

Depending on your application server, you may update the classpath in the administration console, in configuration files or in the startup script for the application server. See documentation for your application server for more information.

4. If the cluster does not have a shared external configuration folder, copy the configuration and extension files from an existing MashZone NextGen Server in the cluster to the new MashZone NextGen Server.

See MashZone NextGen File-Based Configuration and Extensions (page 215) for a list of files and folders to copy.

5. Copy the server configuration that cannot be shared from an existing MashZone NextGen Server in the cluster to the new MashZone NextGen Server. See MashZone NextGen File-Based Configuration and Extensions (page 215) for details on the files and locations for this step.
6. Update the application server that hosts the new MashZone NextGen Server with the same cluster configuration as other cluster members.

In addition to the basic cluster configuration required by your application server and load balancer, MashZone NextGen has a single requirement for application-layer cluster configuration. You must either:

Enable session replication in each application server in the cluster.

Enable session affinity, sometimes also called 'sticky sessions,' in the load balancer.

Or do both.

See documentation for your application server and/or load balancer for information on how to do this.

- Restart the new MashZone NextGen Server.

8.5.6 Sharing the MashZone NextGen Repository in Clustered Environments

In clustered environments, all MashZone NextGen Servers in the cluster must work with a single, shared MashZone NextGen Repository. You can Create and Share a New MashZone NextGen Repository (page 212) with cluster members, typically when you are creating new environments. Or you can Share an Existing MashZone NextGen Repository (page 213) within a cluster.

8.5.6.1 Create and Share a New MashZone NextGen Repository

To create a new shared repository

- Create a new MashZone NextGen Repository in the appropriate database for your environment. See the table below.
- Copy the JAR file for the JDBC driver for your database to the **<MashZone NextGen installation>/apache-tomcat/lib** folder on all cluster nodes.
- If this is a new cluster, update configuration information for the MetaData, User and Snapshot repositories for one MashZone NextGen Server in the cluster. See steps 3 onward in Move the MashZone NextGen repository to a robust database solution (page 116) for instructions.
- Enable distributed caching for artifacts. See Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for more information and instructions.
- If the MashZone NextGen Repository is hosted in Microsoft SQL Server, MySQL or Oracle, change the repository JAR in the MashZone NextGen Server.
- Restart each MashZone NextGen Server in the cluster.

Using the SQL tool for the database that will be host, add MashZone NextGen Repository tables with the scripts shown below from the corresponding folder in MashZoneNG-install/prestorepository:

1. Database	Folder	SQL Scripts
Microsoft SQL Server	mssqlldb	▪ createDBTables.txt for MetaData and the default User Repository

MySQL	mysqlldb	<ul style="list-style-type: none"> ▪ createSnapsTables.txt for Snapshots ▪ createSchedulerTables.txt for Scheduler
Oracle	oracledb	
PostgreSQL	postgresdb	

There are also scripts to drop the corresponding MashZone NextGen Repository tables in these folders, if needed.

8.5.6.2 Share an Existing MashZone NextGen Repository

If you are creating a cluster using an existing MashZone NextGen Repository or simply adding members to an existing cluster, you simply update each new MashZone NextGen Server in the cluster to use the existing repository.

Procedure

1. If the cluster does not have a shared JDBC driver folder and a shared external configuration folder. Copy the JAR file for the JDBC driver for your database to the MashZoneNG-install/apache-tomcat/lib folder for the new MashZone NextGen Server cluster member.
See Setting Up an External MashZone NextGen Configuration Folder (page 213) for more information on shared configuration for clusters.
2. Enable distributed caching for artifacts (required) and optionally distributed caching for mashable/mashup responses. See Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for more information and instructions.
3. If the MashZone NextGen Repository is hosted in Microsoft SQL Server, MySQL or Oracle, change the repository JAR in the MashZone NextGen Server.
4. Restart the new MashZone NextGen Server for this cluster.

8.5.7 Setting Up an External MashZone NextGen Configuration Folder

Most configuration for MashZone NextGen and most of the extensions that you add for your organization's use are stored in the MashZone NextGen Repository. However, some MashZone NextGen configuration and extensions are file based.

By default, MashZone NextGen keeps configuration and extensions in the MashZone NextGen Server web application in these folders:

- <MashZone NextGen installation/apache-tomcat/mashzone/WEB-INF/classes for class, configuration and extension files
- <MashZone NextGen installation/apache-tomcat/mashzone/WEB-INF/lib and
- <MashZone NextGen installation/apache-tomcat/mashzone/WEB-INF/config for JAR files.

You can move most of these configuration and extension files to folders that are external to the MashZone NextGen Server.

Important: MashZone NextGen documentation refers to all of these folders as MashZoneNG-config.

Using external configuration folders for MashZone NextGen is a best practice as they simplify deployment and upgrades of the MashZone NextGen Server. They also simplify configuration management for clustered environments. External configuration folders are not required, however.

To create and use an external configuration folder for MashZone NextGen

- Create the top-level external folder to use for MashZone NextGen configuration, such as **PrestoConfig**. In clustered environments, share or mount this folder across the entire cluster.
- You can create subfolders under this external folder to organize configuration and extensions.
- For clustered environments, create subfolders under the top-level external configuration folder for:
 - The standard **classes** and **lib** folders.
 - Built-in and user-defined functions for use in RAQL queries for MashZone NextGen Analytics. See Configure, Compile, Deploy and Test User-Defined Functions for more information.
 - If not complete, finish configuration for the MashZone NextGen Server and move the configuration and extension files to the external configuration folder or an appropriate subfolder. See the MashZone NextGen File-Based Configuration and Extensions (page 215) section for the specific configuration steps, files and locations.
- Add the external MashZone NextGen configuration folder, **and any subfolder** that contains extensions or JAR files, to the classpath for the application server(s) hosting the MashZone NextGen Server.
- You may update the classpath in configuration files or in the startup script for the application server.
- For Windows environments, for example, you can edit the tomcat-install/bin/setenv.bat file and update the classpath environmental variable to be something like this:
 - **set**
"CLASSPATH=%CLASSPATH%;C:\PrestoConfig;C:\PrestoConfig\classes;C:\PrestoConfig\lib;C:\PrestoConfig\db\jdbc"
- On Linux, Mac OS X or UNIX systems, you would update tomcat-install/bin/setenv.sh to something like this:
 - **CLASSPATH="\$CLASSPATH":/users/PrestoConfig:/users/PrestoConfig/classes:/users/PrestoConfig/lib:users/PrestoConfig/db/jdbc**

8.5.7.1 MashZone NextGen File-Based Configuration and Extensions

Most file-based configuration or extensions involve information that MashZone NextGen needs to connect to the MashZone NextGen Repository or extensions that must be added to the application server's classpath. In clustered environments, you can share extensions and some of this file-based configuration using an external configuration folder. See MashZone NextGen Configuration Files That Can Be External (page 215) and MashZone NextGen Extensions (page 217) for details on resources that can be shared across a cluster.

Some file-based configuration, however, **must** reside in the web application for each MashZone NextGen Server. In clusters, this configuration must be replicated in each cluster member. See MashZone NextGen Configuration Files That Must Be Internal (page 216) for details.

8.5.7.2 MashZone NextGen Configuration Files That Can Be External

File	Description and Configuration	Default Location
dynamiccache.xml	Default configuration information for dynamic In-Memory Stores created by MashZone NextGen Analytics.	MashZoneNG-install/ apache-tomcat/mash zone/WEB-INF/classe s
ehcache.xml	Configuration information for MashZone NextGen caches. This also contains configuration for MashZone NextGen Analytics In-Memory Stores from version 3.6.	
presto.config	Miscellaneous MashZone NextGen properties, including the path to the deployed web app home folder.	
The Terracotta BigMemory license file	The license file for BigMemory, used for MashZone NextGen caches and MashZone NextGen Analytics In-Memory Stores, is a separate license file from the MashZone NextGen license. You can keep the BigMemory license in an external folder shared across the cluster. See Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for required configuration steps to enable a shared license.	

File	Description and Configuration	Default Location
userRespositoryLdap.properties	Connection information for your LDAP Directory. See Integrate Your LDAP Directory with MashZone NextGen (page 126) for details.	

8.5.7.3 MashZone NextGen Configuration Files That Must Be Internal

The file-based configuration that must remain in each MashZone NextGen Server web application resides in the web-apps-home/mashzone/WEB-INF/classes folder.

For upgrades to new MashZone NextGen versions, you can generally copy these configuration files from your existing MashZone NextGen version to the new version. Review the MashZone NextGen Release Notes for changes or new features that may require updates to configuration.

For clustered environments, you **must** copy these configuration files to each cluster member. In most cases, you change configuration once, when you first deploy a MashZone NextGen Server in the cluster. Any subsequent changes to this configuration for one cluster member, however, must be copied to all other cluster members manually, using a scheduled job or using another replication scheme.

File	Description and Configuration
applicationContext-commonServices.xml	You edit configuration in this file if you choose to use distributed response caching for MashZone NextGen. See Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for more information. You may need to update this configuration, as needed, to add additional distributed cache nodes to tune performance.
applicationContext-security.xml	You edit this file initially to enable either SSO authentication or X509 certificate authentication for MashZone NextGen. See Authentication with Single Sign-On Solutions (page 100) for more information.
applicationContext-scheduler.xml	You edit this file when you move the MashZone NextGen Repository from the default Derby database to a robust solution. See Move the MashZone NextGen repository to a robust database solution (page 116) for more information.
log4j.properties	This file is updated automatically when you change logging configuration in the Admin Console. See Configure Logging for the MashZone NextGen Server (page 164) for details. When you change logging for MashZone NextGen Servers in a cluster, only the specific MashZone NextGen Server that the Admin Console is

	connected to is affected. To change logging for the entire cluster, you must update this file and copy it to each cluster member.
userRepositoryApplicationContext.xml	You edit these files when you configure MashZone NextGen to use your LDAP Directory as the user repository. See Integrate Your LDAP Directory with MashZone NextGen (page 126) for details.
userRepositoryApplicationContext-ldap.xml	

8.5.7.4 MashZone NextGen Extensions

Some extensions, such as macros, are registered and reside in the MashZone NextGen Repository. Any of the following file-based extensions can reside in an external folder:

File	Default Location
Custom filter classes for single sign-on authentication. See Implementing a Custom SSO Filter (page 103) for details.	<MashZone NextGen installation>/apache-tomcat/mashzone/WEB-INF/classes
Classes and third-party libraries for a user-defined function library to use with RAQL. See Create and Add User-Defined Functions for RAQL Queries for more information.	or <MashZone NextGen installation>/apache-tomcat/mashzone/WEB-INF/lib (for JARs)

8.5.8 MashZone NextGen dashboards in a clustered scenario

You can use MashZone NextGen dashboards in a clustered scenario.

The following chapters describe how to configure MashZone NextGen to use dashboards and data feeds in a multiple master-client scenario.

8.5.8.1 Preliminary

Before you can configure MashZone NextGen using in a clustered scenario you have to perform the following steps.

Procedure

1. Install at least two regular MashZone NextGen instances on two different machines.

Software AG Installer enables you to install MashZone NextGen. Detailed information on how to use Software AG Installer is available in the documentation **Using the Software AG Installer**.

2. Connect all instances to the same central database according to section Move the MashZone NextGen repository to a robust database solution (page 116).
3. If MashZone NextGen should consume event-based data, make sure that MashZone NextGen Event Service is started on at least one machine. To run Event Service in high-availability mode, you must run it on at least three different machines.

The preliminary for configuring MashZone NextGen are completed.

8.5.8.2 Configuration

The following chapters describe the relevant configurations of MashZone NextGen Dashboard in a clustered scenario.

From MashZone NextGen 10.2 on, Event Service can be run in clustered high-availability mode. There is no distinction between MashZone NextGen master and slave servers anymore. This means, to achieve high availability, you can now deploy multiple instances of Event Service to form a cluster. Cluster management is implemented using Apache Zookeeper.

Each MashZone NextGen server may connect to an arbitrary Event Service instance in the cluster and receive data from it. In case of an Event Service node failure, MashZone NextGen will automatically and transparently fail over to another alive node. When an Event Service node re-joins the cluster after a node failure, it will restore its state from Event Service' internal HA-Store.

8.5.8.2.1 MashZone NextGen Event Service

It is not recommended to configure Event Service clustering manually. The service contains of three sub-components (Event Service Core, HA-Store, Zookeeper) and the configuration of those components must be coherent. The only officially supported way to configure Event Service is using Command Central plugin.

For illustration purposes, we assume that the Event Service cluster should contain of three machines: machine1, machine2, and machine 3.

Configure the following parameters on your MashZone NextGen Event Service nodes.

8.5.8.3 `rtbs.properties`

Edit the **rtbs.properties** file and set **rtbs.id**, **rtbs.zookeeper.urls**, and **rtbs.replication.factor** as described below.

The **rtbs.properties** file is located in the <MashZone NextGen installation> /rtbs/conf/ directory

- `rtbs.id=<Unique id per node between 1 and 255>`, for example, 1 for machine 1, 2 for machine 2, 3 for machine 3
- `rtbs.zookeeper.urls=machine1:12181,machine2:12181,machine3:12181`

Replace the host names **machine1**, **machine2**, and **machine3** by your real servers' host names and also set the right ports. You must enumerate all Event Service cluster members.

- **`rtbs.replication.factor=<number>`**

`<number>` in replication factor determines how many replicas of persisted data (for high availability) should exist. The number must be `>= 1` and `<=` the number of Event Service instances in the cluster. The number must be `>= 2` to achieve any fault tolerance. As a rule of thumb: the higher the number of replicas the more Event Service node failures can be tolerated. However, the amount of storage required on the Event Service machines increases linearly with the number of replicas.

Note: The replication factor setting only affects buffers that are created after the replication factor was changed, that is, changing the replication factor does not affect existing buffers.

8.5.8.4 `zookeeper.properties`

Edit the **`zookeeper.properties`** file and add the parameters described below.

The **`zookeeper.properties`** file is located in the `<MashZone NextGen installation>/rtbs/zookeeper/conf/` directory

Add the parameters:

- `initLimit=5`
- `syncLimit=2`

The **`initLimit`** parameter is the time in seconds Zookeeper waits for other nodes to join the cluster (Zookeeper Ensemble). If the cluster initialization fails (especially relevant if nodes are not started in an automated and coordinated way), increase the value, for example, to 60.

For each Event Service server in the cluster, add a line as follows:

- **`server.<number>=<hostname>:2888:3888`**

Here, `<number>` is a strictly monotonic increasing number starting at 1. It is not necessarily equivalent to `rtbs.id`, that is, if you decide to use `rtbs.id=42` for machine1, the entry would still be `server.1=machine1:2888:3888`

For our example, this means that we need to add the following three lines:

- `server.1=machine1:2888:3888`
- `server.2=machine2:2888:3888`
- `server.3=machine3:2888:3888`

Create the file `<MashZone NextGen installation>/rtbs/zookeeper/data/zk-snapshot/myid`.

Add the machine's `rtbs.id` to the file, that is, just the digit 1 for machine 1, digit 2 on machine 2, digit 3 for machine 3.

8.5.8.5 server.properties

Edit the file **server.properties**, adjust the **broker.id** parameter to the machine's Event Service ID and adjust the **zookeeper.connect** parameter.

The **server.properties** file is located in the **<MashZone NextGen installation>/rtbs/ha-store/conf/** directory.

- **broker.id**=<rtbs.id>, that is 1 on machine 1, 2 on machine 2, 3 on machine 3
- **zookeeper.connect**=machine1:12181,machine2:12181,machine3:12181

8.5.8.6 MashZone NextGen nodes

Configure the following parameters on your MashZone NextGen nodes.

Edit the **mashzone.rtbs.url** parameter in the **mashzone.properties** file and set host name and port accordingly.

<MashZone NextGen

installation>/apache-tomcat/webapps/mashzone/WEB-INF/mashzone.properties

- **mashzone.rtbs.url**=**machine1:12181,machine2:12181,machine3:12181**

Replace the host names **machine1**, **machine2**, and **machine3** by your real servers' host names and also set the right ports. You must enumerate all Event Service cluster members.

Note: Calls to the Event Service API are server-to-server calls which usually happen behind the load balancer. In some cases it might be necessary to route these calls through the load balancer as well. In this case, make sure that you configure your load balancer accordingly. Set the load balancer's host name and port in the parameters mentioned above.

8.5.8.6.1 Customizing dashboards

MashZone NextGen dashboards can be customized by adding custom style templates for the dashboard application and the dashboard content. Additionally custom widgets can be created via the pluggable widget framework. If these options shall be applied in a clustered scenario, you must synchronize the relevant folders and restart MashZone NextGen on all nodes of the cluster.

8.5.8.7 Custom styles

By default, custom style templates available are stored in the following folders.

- <MashZone NextGen
installation>/apache-tomcat/webapps/mashzone/hub/dashboard/assets/custom-look-and-feel/application

- <MashZone NextGen installation>/apache-tomcat/webapps/mashzone/hub/dashboard/assets/custom-look-and-feel/dashboard

To apply the custom templates on all cluster nodes, make sure that these folders are synchronized on all machines. Since the less files need to be compiled before the styles can be used, MashZone NextGen has to be restarted on all cluster nodes.

8.5.8.8 Custom widgets

By default, custom widgets available are stored in the following folders.

<MashZone NextGen installation>/apache-tomcat/webapps/mashzone/hub/dashboard/widgets/customWidgets

To make the custom widgets available on all cluster nodes, make sure that the folders is synchronized on all machines. In this case, restarting MashZone NextGen on all cluster nodes is required as well.

8.5.8.9 Using JDBC drivers

JDBC driver binaries have to be available on every cluster node to allow class loading in the JVM. Since MashZone NextGen version 9.10 the binaries are stored in the DB and restored in <MashZone NextGen installation>\apache-tomcat\webapps\mashzone\WEB-INF\config\db\jdbc on all cluster nodes if not available. Automatic class loading on demand works fine, so that no further steps have to be taken to make JDBC resources available in a clustered scenario.

8.5.8.10 Local file resources

Local file resources are not recommended and not supported in a clustered scenario due to synchronization issues.

In a Windows landscape it might be possible to use such file resources by mapping the same network drive to the same network share. There may also be other file sharing mechanisms working in other OS landscapes, but URL based access is preferable.

8.6 Event Service Configuration and Administration

8.6.1 About the Event Service and Event Data

The Event Service allows MashZone NextGen to connect to the Event Bus for Software AG and subscribe to events published by other Software AG applications.

The Event Bus handles events published by **event producers** which may be a variety of Software AG applications. It routes events as they are published to **event consumers**, such as MashZone NextGen, who have subscribed to specific **event types**.

MashZone NextGen also uses the Event Service to connect to Apama and to subscribe to and work with events from **scenarios** (also sometimes called **dataviews**). Apama scenarios have event data that has been specifically transformed for use in dashboards.

For MashZone NextGen, each Event Bus or Apama subscription feeds an **event source** which is managed by the Event Service. Event sources receive events for subscriptions, store them in memory and act as the data source for the corresponding EDA, DES, or Apama event.

From MashZone NextGen 10.2 on, Event Service provides high-availability mode and horizontal scalability. It can be run in a cluster to achieve both tolerance against failure of cluster members and horizontal scalability by distributing requests from MashZone NextGen to all cluster members.

Cluster members may be restarted (for example, after a fault or maintenance reboot) and automatically re-join the cluster and restore their state from disk. See MashZone NextGen dashboards in a clustered scenario (page 217), MashZone NextGen Event Service (page 218) and Command Central plug-in (page 137) for details concerning cluster configuration.

8.6.2 Use Events as Information Sources

To use events as information sources, you must create EDA, DES, or Apama event sources. Unlike other sources, only MashZone NextGen administrators can create EDA, DES and Apama event sources for other users to work with.

To create these sources, administrators must create event sources as needed. MashZone NextGen retrieves event data from these event sources when users run the corresponding EDA, DES or Apama event source. If views for the event source are real-time views, events are pushed to the view automatically.

- Create EDA Event Sources (Seite 223)
- Create DES Event Source (Seite 244)
- Create Apama Event Sources (Seite 233)
- Start or Stop an Event Source (Seite 254)

You can also Manage Apama Event Sources (page 232), Manage DES Event Sources (Seite 244), and Manage EDA Event Sources (page 222).

8.6.3 Manage EDA Event Sources

To identify, create, edit, delete, import or export **EDA Event Sources**:

1. In the program bar click the user name by which you are logged in to MashZone NextGen.
2. Click Admin Console.
3. Click Event Service to expand this section of the Administration menu.

4. Click Event Service. The Event Service page will be displayed.
5. Open the EDA tab.
6. Select further steps:
 - Create EDA Event Sources (Seite 223)
 - Edit EDA Event Sources (Seite 227)
 - Duplicate EDA Event Sources (Seite 231)
 - Delete EDA Event Sources (Seite 231)
 - Share EDA Event Sources (Seite 232)

8.6.3.1 Create EDA Event Sources

You can register subscriptions with the Event Bus. This creates **EDA Event Sources** that hold published events in memory and a corresponding event mashable in MashZone NextGen.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **EDA** tab.
5. Click **Create EDA Event Source**.
6. Set the properties for this event source. See table **EDA Event Source properties** below.
7. Click **Save**.

The EDA Event Source is created and listed by alias name.

EDA Event Source properties

Property	Required	Description
Alias	yes	Enter a unique name for this event source.
Start event source automatically on server startup		This option is set by default, which automatically starts this event source when MashZone NextGen Server starts. Clear this option if you need to manually control startup for this event source.
Event type	yes	Click  Refresh to update the list of event types. Select the type of the event this event source should subscribe to. The XML schema files for these event types must exist in the Event Type Store directory of MashZone NextGen Event Service. The Event Type Store directory can be configured using Command Central.
Filter predicates		Enter a filter expression defining the events to be published to this event source.

Check validity		<p>Available only if Strategy is set to Buffer.</p> <p>Determines whether saved events are valid with respect to the current time frame for the application (ta) and removes invalid events from the event source.</p> <ul style="list-style-type: none"> ▪ An event has a time stamp in the form of a time interval (I) = Start time - End time [ts - te); with ts being an element of I, and te not being an element of I. ▪ The current time of the application is determined by the start time of the last received event. ▪ An event is valid if the current time of the application is within the interval, i.e., [ts <= ta < te).
Preprocess and filter heartbeats		<p>Removes empty events with no data from the event source. Empty events can, however, update the application time and thus can force a consolidation of the event source content.</p>
Strategy	yes	<p>The strategy that this event source uses for saving and removing events published from the Event Bus. Valid strategies are:</p> <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until event source memory reaches capacity and then the event source removes the oldest events. ▪ DELTA= Events are stored by ID and added, updated or removed based on a command within the event. An event with an Insert command is saved in event source memory, any existing event with the same ID is overwritten. An event with a Remove command removes an existing event with the same ID.
Consider dimension		<p>Available only when Strategy is set to BUFFER.</p> <p>Set this option to save events in separate series (or buckets) for each unique value of the Dimension attribute.</p>
Dimension attribute	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Select the event attribute whose unique values determine separate event series (buckets) for this event source.</p>
Max. number of dimension values	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of unique dimension values (buckets) that can be tracked. Thus this is the maximum number of series that can store events.</p> <p>Default value: 1 (Max: 100.000).</p> <p>The product of Max. number of dimension values and</p>

		<p>Capacity per dimension value must not be greater than 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Determines how additional events are handled if they have new unique values for the dimension that defines buckets in this event source but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached.</p> <p>Set this option to change the bucket strategy to FIFO (first-in, first-out) which discards events for older series (buckets) and stores the newer event in a new series (bucket).</p> <p>Default value: false</p>
Capacity per dimension value	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value.</p> <p>Default value: 10</p> <p>The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.</p>
Event ID attribute	yes	<p>Available only when Strategy is set to DELTA.</p> <p>Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.</p>
Command attribute	yes	<p>Available only when Strategy if set to DELTA.</p> <p>Select the attribute that contains the event command (Insert or Remove). The event ID and command determines which events are stored, updated or removed in this event source.</p>
Capacity	yes	<p>Enter the maximum number of events to store in this event source. (Max: 100.000)</p> <p>Default value: 10</p>
Memory model		<p>Determines where events are stored:</p> <ul style="list-style-type: none"> ▪ Internal: the default which stores events in local memory for this event source. ▪ BigMemory: stores events in a local BigMemory cache.
Throttling		<p>Controls the speed and volume of event data that is pushed to</p>

	<p>views that subscribe to this event source. By default, event sources push event data every 500 milliseconds. You can:</p> <ul style="list-style-type: none"> ▪ Change the number of milliseconds to control throttling. ▪ Change the measurement (Default=500) to Events to have throttling wait until a specific number of events are received and change the number, if needed. <p>See example below.</p>
Exception	<p>Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1):</p> <ul style="list-style-type: none"> ▪ A number ▪ Milliseconds or Events as the measurement for the exception criteria <p>See example below</p>

Simple and Hybrid Throttling Strategies

Simple throttling strategies cause an event source to wait for either a specific time interval or for the receipt of a specific number of events and then push all new events to any subscribing real-time views. Throttling can slow event updates to real-time views when the volume or frequency for events causes rendering issues.

The default behavior is to push events to views based on a time interval of every 500 milliseconds. You can change the time interval or change the criteria to push events once a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies may still not even out event flow adequately. Instead, you can create hybrid strategies, such as "generally push every 50 milliseconds, but at most 10 events."

Hybrid strategies define the general throttling with the **Throttling** fields. You set the **Exception** option and define the exception that should break the general rule in the **Exception** criteria fields:

Throttling * Exception...

With the example hybrid throttling strategy shown above:

- The event source would wait 50 milliseconds after pushing events to subscribing views.
- If less than 10 events are received in that 50 milliseconds, they are pushed at the end of the interval.
- If a tenth event is received within the 50 milliseconds, these 10 events are pushed to subscribing views and both the time interval and the count of events begins again.
- If no events are received within the time interval, the event source waits until it receives an event. When an event is received, the event source pushes this event to subscribing views and restarts the time interval.

On the EDA Event Source overview page you can click on the **Alias** to show a preview of the specific Event Source properties.

8.6.3.2 Edit EDA Event Sources

You can edit already existing EDA Event Source.

Changes in EDA connection properties can immediately affect data feed calculations so that they may not execute properly.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **EDA** tab.
5. Click the  **Edit** icon to configure a specific **EDA** connection.
6. Set the properties for this event source:

Table 1. EDA Event Source properties

Property	Required	Description
Alias	yes	Enter a unique name for this event source.
Start event source automatically on server startup		This option is set by default, which automatically starts this event source when MashZone NextGen Server starts. Clear this option if you need to manually control startup for this event source.
Event type	yes	Click  Refresh to update the list of event types. Select the type of the event this event source should subscribe to. The XML schema files for these event types must exist in the Event Type Store directory of MashZone NextGen Event Service. The Event Type Store

		directory can be configured using Command Central.
Filter predicates		Enter a filter expression defining the events to be published to this event source.
Check validity		<p>Available only if Strategy is set to Buffer.</p> <p>Determines whether saved events are valid with respect to the current time frame for the application (ta) and removes invalid events from the event source.</p> <ul style="list-style-type: none"> ▪ An event has a time stamp in the form of a time interval (I) = Start time - End time [ts - te); with ts being an element of I, and te not being an element of I. ▪ The current time of the application is determined by the start time of the last received event. ▪ An event is valid if the current time of the application is within the interval, i.e., [ts <= ta < te).
Preprocess and filter heartbeats		Removes empty events with no data from the event source. Empty events can, however, update the application time and thus can force a consolidation of the event source content.
Strategy	yes	<p>The strategy that this event source uses for saving and removing events published from the Event Bus. Valid strategies are:</p> <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until event source memory reaches capacity and then the event source removes the oldest events. ▪ DELTA= Events are stored by ID and added, updated or removed based on a command within the event. An event with an Insert command is saved in event source memory, any existing event with the same ID is overwritten. An event with a Remove command removes an existing event with the same ID.
Consider dimension		<p>Available only when Strategy is set to BUFFER.</p> <p>Set this option to save events in separate series (or buckets) for each unique value of the Dimension attribute.</p>
Dimension attribute	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Select the event attribute whose unique values determine separate event series (buckets) for this event source.</p>
Max. number of dimension values	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of unique dimension values (buckets) that can be tracked. Thus this is the maximum number of series that</p>

		<p>can store events.</p> <p>Default value: 1 (Max: 100.000).</p> <p>The product of Max. number of dimension values and Capacity per dimension value must not be greater than 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Determines how additional events are handled if they have new unique values for the dimension that defines buckets in this event source but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached.</p> <p>Set this option to change the bucket strategy to FIFO (first-in, first-out) which discards events for older series (buckets) and stores the newer event in a new series (bucket).</p> <p>Default value: false</p>
Capacity per dimension value	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value.</p> <p>Default value: 10</p> <p>The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.</p>
Event ID attribute	yes	<p>Available only when Strategy is set to DELTA.</p> <p>Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.</p>
Command attribute	yes	<p>Available only when Strategy if set to DELTA.</p> <p>Select the attribute that contains the event command (Insert or Remove). The event ID and command determines which events are stored, updated or removed in this event source.</p>
Capacity	yes	<p>Enter the maximum number of events to store in this event source. (Max: 100.000)</p> <p>Default value: 10</p>
Memory model		<p>Determines where events are stored:</p> <ul style="list-style-type: none"> ▪ Internal: the default which stores events in local memory for this event source. ▪ BigMemory: stores events in a local BigMemory cache.

Throttling		<p>Controls the speed and volume of event data that is pushed to views that subscribe to this event source. By default, event sources push event data every 500 milliseconds. You can:</p> <ul style="list-style-type: none"> ▪ Change the number of milliseconds to control throttling. ▪ Change the measurement (Default=500) to Events to have throttling wait until a specific number of events are received and change the number, if needed.
Exception		<p>Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1):</p> <ul style="list-style-type: none"> ▪ A number ▪ Milliseconds or Events as the measurement for the exception criteria

8.6.3.2.1 Simple and Hybrid Throttling Strategies

Simple throttling strategies cause an event source to wait for either a specific time interval or for the receipt of a specific number of events and then push all new events to any subscribing real-time views. Throttling can slow event updates to real-time views when the volume or frequency for events causes rendering issues.

The default behavior is to push events to views based on a time interval of every 500 milliseconds. You can change the time interval or change the criteria to push events once a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies may still not even out event flow adequately. Instead, you can create hybrid strategies, such as "generally push every 50 milliseconds, but at most 10 events."

Hybrid strategies define the general throttling with the **Throttling** fields. You set the **Exception** option and define the exception that should break the general rule in the **Exception** criteria fields:

Throttling * Exception...

With the example hybrid throttling strategy shown above:

- The event source would wait 50 milliseconds after pushing events to subscribing views.
- If less than 10 events are received in that 50 milliseconds, they are pushed at the end of the interval.
- If a tenth event is received within the 50 milliseconds, these 10 events are pushed to subscribing views and both the time interval and the count of events begins again.
- If no events are received within the time interval, the event source waits until it receives an event. When an event is received, the event source pushes this event to subscribing views and restarts the time interval.

Related tasks

Create EDA Event Sources (Seite 223)

8.6.3.3 Duplicate EDA Event Sources

You can duplicate existing EDA Event Source.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page will be displayed.
4. Open the **EDA** tab.
5. Click the  **Copy** icon to duplicate a specific EDA Event Source.

The selected EDA Event Source is duplicated and listed with the prefix **copy_** in the **Alias**.

8.6.3.4 Delete EDA Event Sources

You can delete existing EDA Event Source.

Warning

Deleting EDA Event Sources may cause data feeds to fail.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.

3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **EDA** tab.
5. Click the  **Delete** icon to delete a specific EDA Event Source.

The selected EDA Event Source is deleted from the list.

8.6.3.5 Share EDA Event Sources

You can share EDA Event Sources with particular users and user groups so that these have access to **EDA Event Services**.

Prerequisite

You have administration privileges.

Regardless of the share, users with administration privilege can access all EDA Event sources.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **EDA** tab.
5. Click the  **Edit event service permissions** icon of the EDA Event Service you want to share.
6. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
7. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
8. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

By default, the owner of the **EDA Event Services** is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.

9. Activate or deactivate the **Display** or **Usage** privileges of a user or user group.

A user or user group with **Display** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Usage** privilege has access to the relevant alias in the data source operator.

10. Click **OK**.

Your changes are applied.

8.6.4 Manage Apama Event Sources

You can create, edit, delete, import, and share **Apama** Event Sources:

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **Apama** tab.
5. Select further steps:
 - Create Apama Event Sources (Seite 233)
 - Edit Apama Event Sources (Seite 238)
 - Duplicate Apama Event Sources (page 238)
 - Delete Apama Event Sources (Seite 238)
 - Share Apama Event Sources (Seite 243)

8.6.4.1 Create Apama Event Sources

MashZone NextGen can work with events published from Apama through the Event Bus. In many cases, however, the events and data you need are defined in Apama **scenarios** which are not accessible through the Event Bus.

To work with Apama scenario events, you must create an Apama Event Source to receive scenario events.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click **Create Apama Event Source**.
6. Set the properties for this event source. See table below.
7. Click **Save**.

The **Apama Event Source** is created and listed by alias name.

Apama Event Source properties

Property	Required	Description
Alias	yes	Enter a unique name for this event source.
Start event source automatically on server startup		This option is set by default, which automatically starts this event source when MashZone NextGen Server starts. Clear this option if you need to manually control startup for this event source.
Apama instance	yes	Alias with the pre-configured connection specification of a running Apama system (local or remote). See Manage

		Apama Instances (Seite 255) for details.
Apama Scenario	yes	Click  Refresh to update the list of Apama scenarios for the selected Apama Event Source. If the Apama URL is set to a valid Apama system, then it is possible to select a scenario this event source should subscribe to.
Strategy	yes	<p>The strategy that this event source uses for saving and removing events published from the Event Bus. Valid strategies are:</p> <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until event source memory reaches capacity and then the event source removes the oldest events. ▪ DELTA= Events are stored by ID and added, updated or removed based on a command within the event. An event with an Insert command is saved in event source memory, any existing event with the same ID is overwritten. An event with a Remove command removes an existing event with the same ID. ▪ PARTIAL_EVENTS = Each event has a unique identifier defined by one or more key fields (see Key attributes). Events contain additional fields, but may not contain all fields possible for the event. Simply put, each event may contain partial data. <p>The event source maintains a single row for each unique event key representing the current full status for that event. Events published by Apama scenarios update the fields in that event source row that are included in the event, leaving other existing data for that event source row intact.</p> <p>Events that have a new unique key are saved as a new row until event source memory is full. Once the event source memory is full, new events are discarded.</p> <ul style="list-style-type: none"> ▪ SCENARIO = Events are stored according to the respective Apama Scenario Definition. <p>MashZone NextGen automatically determines the unique identifier of each event, defined by one or more key fields. It maintains a single row for each event key, updating its values by processing the Apama scenario</p> <p>InstanceUpdated Events.</p> <p>This strategy also supports the Apama scenario InstanceRemoved Events. When an InstanceRemoved event is received, the</p>

		<p>corresponding row is removed from the Apama Event Source.</p> <p>Apama Event Sources using the SCENARIO strategy do not store events in the Event Service HA-Store. When a SCENARIO Event Source is restarted, it always receives the current scenario state from its Apama instance connection.</p>
Consider dimension		<p>Available only when Strategy is set to BUFFER.</p> <p>Set this option to save events in separate series (or buckets) for each unique value of the Dimension attribute.</p>
Dimension attribute	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Select the event attribute whose unique values determine separate event series (buckets) for this event source.</p>
Max. number of dimension values	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of unique dimension values (buckets) that can be tracked. Thus this is the maximum number of series that can store events.</p> <p>Default value: 10 (Max: 100.000).</p> <p>The product of Max. number of dimension values and Capacity per dimension value must not be greater than 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Determines how additional events are handled if they have new unique values for the dimension that defines buckets in this event source but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached.</p> <p>Set this option to change the bucket strategy to FIFO (first-in, first-out) which discards events for older series (buckets) and stores the newer event in a new series (bucket).</p> <p>Default value: false</p>
Capacity per dimension value	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p>

		<p>Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value.</p> <p>Default value: 10;</p> <p>The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.</p>
Event ID attribute	yes	<p>Available only when Strategy is set to DELTA.</p> <p>Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.</p>
Command attribute	yes	<p>Available only when Strategy if set to DELTA.</p> <p>Select the attribute that contains the event command (Insert or Remove). The event ID and command determines which events are stored, updated or removed in this event source.</p>
Key attributes	yes	<p>Available only when Strategy is set to PARTIAL_EVENT.</p> <p>The field(s) in events with partial data that uniquely identify an event. The event ID is used to ensure that events with partial data properly insert or update events in this event source.</p> <p>Select one or more attributes that uniquely identify events for this Apama scenario. If multiple fields are required, the order in which you select attributes determines how fields are combined to determine event IDs.</p>
Capacity	yes	<p>Enter the maximum number of events to store in this event source. (Max: 100.000)</p> <p>Default value: 10</p>
Memory model		<p>Determines where events are stored:</p> <ul style="list-style-type: none"> ▪ Internal: the default which stores events in local memory for this event source. ▪ BigMemory: stores events in a local BigMemory cache.
Throttling		<p>Controls the speed and volume of event data that is pushed to views that subscribe to this event source. By default, event sources push event data every 500 milliseconds. You can:</p> <ul style="list-style-type: none"> ▪ Change the number of milliseconds to control throttling.

	<ul style="list-style-type: none"> Change the measurement (Default=500) to Events to have throttling wait until a specific number of events are received and change the number, if needed. <p>See Example below.</p>
Exception	<p>Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1):</p> <ul style="list-style-type: none"> A number Milliseconds or Events as the measurement for the exception criteria <p>See Example below.</p>

Simple and Hybrid Throttling Strategies

Simple throttling strategies cause an event source to wait for either a specific time interval or for the receipt of a specific number of events and then push all new events to any subscribing real-time views. Throttling can slow event updates to real-time views when the volume or frequency for events causes rendering issues.

The default behavior is to push events to views based on a time interval of every 500 milliseconds. You can change the time interval or change the criteria to push events once a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies may still not even out event flow adequately. Instead, you can create hybrid strategies, such as "generally push every 50 milliseconds, but at most 10 events."

Hybrid strategies define the general throttling with the **Throttling** fields. You set the **Exception** option and define the exception that should break the general rule in the **Exception** criteria fields:

Throttling * Exception...

With the example hybrid throttling strategy shown above:

- The event source would wait 50 milliseconds after pushing events to subscribing views.

- If less than 10 events are received in that 50 milliseconds, they are pushed at the end of the interval.
- If a tenth event is received within the 50 milliseconds, these 10 events are pushed to subscribing views and both the time interval and the count of events begins again.
- If no events are received within the time interval, the event source waits until it receives an event. When an event is received, the event source pushes this event to subscribing views and restarts the time interval.

On the Apama Event Source overview page you can click on the **Alias** to show a preview of the specific Event Source properties.

8.6.4.2 Duplicate Apama Event Sources

You can duplicate an existing Apama Event Source.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Copy** icon to duplicate a specific Apama Event Source.

The selected Apama Event Source is duplicated and listed with the prefix **copy_** in the **Alias**.

8.6.4.3 Delete Apama Event Sources

You can delete an Apama Event Source.

Note: Deleting an Apama Event Source may cause data feeds to fail.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Delete** icon to delete a specific Apama Event Source.
6. Click **Save**.

The selected Apama Event Source is deleted from the list.

8.6.4.4 Edit Apama Event Sources

You can edit an already existing Apama Event Source.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Edit** icon to configure an Apama Event Source.
6. Set the properties for this event source. See table **Apama Event Source properties** below.
7. Click **Save**.

The **Apama Event Source** is created and listed by alias name.

Apama Event Source properties

Property	Required	Description
Alias	yes	1. Enter a unique name for this event source.
Start event source automatically on server startup		This option is set by default, which automatically starts this event source when MashZone NextGen Server starts. Clear this option if you need to manually control startup for this event source.
Apama instance	yes	Alias with the pre-configured connection specification of a running Apama system (local or remote). See Manage Apama Instances (Seite 255) for details.
Apama Scenario	yes	Click  Refresh to update the list of Apama scenarios for the selected Apama Event Source. If the Apama URL is set to a valid Apama system, then it is possible to select a scenario this event source should subscribe to.
Strategy	yes	The strategy that this event source uses for saving and removing events published from the Event Bus. Valid strategies are: <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until event source memory reaches capacity and then the event source removes the oldest events. ▪ DELTA= Events are stored by ID and added, updated or removed based on a command within the event. An event with an Insert command is saved in event source memory, any existing event with the same ID is overwritten. An event with a Remove command removes an existing event with the same ID. ▪ PARTIAL_EVENTS = Each event has a unique identifier defined by one or more key fields (see Key attributes). Events contain additional fields, but may not contain all fields possible for the event. Simply put, each event may contain partial data.

		<p>The event source maintains a single row for each unique event key representing the current full status for that event. Events published by Apama scenarios update the fields in that event source row that are included in the event, leaving other existing data for that event source row intact.</p> <p>Events that have a new unique key are saved as a new row until event source memory is full. Once the event source memory is full, new events are discarded.</p> <ul style="list-style-type: none"> ▪ SCENARIO = Events are stored according to the respective Apama Scenario Definition. <p>MashZone NextGen automatically determines the unique identifier of each event, defined by one or more key fields. It maintains a single row for each event key, updating its values by processing the Apama scenario InstanceUpdated Events.</p> <p>This strategy also supports the Apama scenario InstanceRemoved Events. When an InstanceRemoved event is received, the corresponding row is removed from the Apama Event Source.</p> <p>Apama Event Sources using the SCENARIO strategy do not store events in the Event Service HA-Store. When a SCENARIO Event Source is restarted, it always receives the current scenario state from its Apama instance connection.</p>
Consider dimension		<p>Available only when Strategy is set to BUFFER.</p> <p>Set this option to save events in separate series (or buckets) for each unique value of the Dimension attribute.</p>
Dimension attribute	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Select the event attribute whose unique values determine separate event series (buckets) for this event source.</p>
Max. number of dimension values	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of unique dimension values (buckets) that can be tracked. Thus this is the maximum number of series that can store events.</p> <p>Default value: 10 (Max: 100.000).</p> <p>The product of Max. number of dimension values and Capacity per dimension value must not be greater than 100.000.</p>
Dimension		<p>Available only when Strategy is set to BUFFER and required</p>

Squeeze-out		<p>when the Consider dimension option is set.</p> <p>Determines how additional events are handled if they have new unique values for the dimension that defines buckets in this event source but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached.</p> <p>Set this option to change the bucket strategy to FIFO (first-in, first-out) which discards events for older series (buckets) and stores the newer event in a new series (bucket).</p> <p>Default value: false</p>
Capacity per dimension value	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value.</p> <p>Default value: 10;</p> <p>The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.</p>
Event ID attribute	yes	<p>Available only when Strategy is set to DELTA.</p> <p>Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.</p>
Command attribute	yes	<p>Available only when Strategy if set to DELTA.</p> <p>Select the attribute that contains the event command (Insert or Remove). The event ID and command determines which events are stored, updated or removed in this event source.</p>
Key attributes	yes	<p>Available only when Strategy is set to PARTIAL_EVENT.</p> <p>The field(s) in events with partial data that uniquely identify an event. The event ID is used to ensure that events with partial data properly insert or update events in this event source.</p> <p>Select one or more attributes that uniquely identify events for this Apama scenario. If multiple fields are required, the order in which you select attributes determines how fields are combined to determine event IDs.</p>
Capacity	yes	<p>Enter the maximum number of events to store in this event source. (Max: 100.000)</p> <p>Default value: 10</p>
Memory model		Determines where events are stored:

		<ul style="list-style-type: none"> ▪ Internal: the default which stores events in local memory for this event source. ▪ BigMemory: stores events in a local BigMemory cache.
Throttling		<p>Controls the speed and volume of event data that is pushed to views that subscribe to this event source. By default, event sources push event data every 500 milliseconds. You can:</p> <ul style="list-style-type: none"> ▪ Change the number of milliseconds to control throttling. ▪ Change the measurement (Default=500) to Events to have throttling wait until a specific number of events are received and change the number, if needed. <p>See Example below.</p>
Exception		<p>Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1):</p> <ul style="list-style-type: none"> ▪ A number ▪ Milliseconds or Events as the measurement for the exception criteria <p>See Example below.</p>

Simple and Hybrid Throttling Strategies

Simple throttling strategies cause an event source to wait for either a specific time interval or for the receipt of a specific number of events and then push all new events to any subscribing real-time views. Throttling can slow event updates to real-time views when the volume or frequency for events causes rendering issues.

The default behavior is to push events to views based on a time interval of every 500 milliseconds. You can change the time interval or change the criteria to push events once a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies may still not even out event flow adequately. Instead, you can create hybrid strategies, such as "generally push every 50 milliseconds, but at most 10 events."

Hybrid strategies define the general throttling with the **Throttling** fields. You set the **Exception** option and define the exception that should break the general rule in the **Exception** criteria fields:

Throttling * Exception...

With the example hybrid throttling strategy shown above:

- The event source would wait 50 milliseconds after pushing events to subscribing views.
- If less than 10 events are received in that 50 milliseconds, they are pushed at the end of the interval.
- If a tenth event is received within the 50 milliseconds, these 10 events are pushed to subscribing views and both the time interval and the count of events begins again.
- If no events are received within the time interval, the event source waits until it receives an event. When an event is received, the event source pushes this event to subscribing views and restarts the time interval.

On the Apama Event Source overview page you can click on the **Alias** to show a preview of the specific Event Source properties.

8.6.4.5 Share Apama Event Sources

You can share Apama Event Sources with particular users and user groups so that these have access to **Apama Event Services**.

Prerequisite

You have administration privileges.

Regardless of the share, users with administration privilege can access all Apama Event sources.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Edit event service permissions** icon of the Apama Event Service you want to share.
6. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
7. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
8. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

By default, the owner of the **Apama Event Services** is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.

9. Activate or deactivate the **Display** or **Usage** privileges of a user or user group.

A user or user group with **Display** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Usage** privilege has access to the relevant alias in the data source operator.

10. Click **OK**.

Your changes are applied.

8.6.5 Manage DES Event Sources

You can manage your DES Event Sources in the **Admin console**.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **Digital Event Services** tab.
5. Select further steps:

Create DES Event Source (Seite 244)

Edit DES Event Sources (Seite 249)

Duplicate DES Event Sources (Seite 253)

Delete DES Event Sources (Seite 253)

Share DES Event Sources (Seite 253)

Create DES Event Source (Seite 244)

Edit DES Event Sources (Seite 249)

Duplicate DES Event Sources (Seite 253)

Delete DES Event Sources (Seite 253)

Share DES Event Sources (Seite 253)

You can share DES Event Sources with particular users and user groups so that these have access to **DES Event Services**.

Activate DES in MashZone NextGen (Seite 254)

To use Digital Event Services (DES), a valid DES license file must be present in the MashZone NextGen installation.

8.6.5.1 Create DES Event Source

You can register subscriptions with the Event Bus. This creates **DES Event Services** that hold published events in memory and a corresponding event mashable in MashZone NextGen.

Starting with MashZone NextGen 10.2, Event Service subscribes using a durable subscription. The durable subscription remains active until the Event Source is deleted, that is, Universal Messaging will buffer events until MashZone NextGen Event Service consumes them. This ensures that no event data is lost in situations where MashZone NextGen Event Service is unavailable. Event Service uses the defined alias name (see below) as the subscription’s identifier. Thus, you need to make sure that there must not be configured two independent Event Service installations that access the same Universal Messaging Channel using the same alias name. That is, if you operate two Event Service Clusters (single or multi-node) and they share a single Universal Messaging Instance, make sure to use different alias names for the same channels.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **Digital Event Services** tab.
5. Click **Create DES Event Source**.
6. Set the properties for this event source. See table below.
7. Click **Save**.

The DES Event Service is created and listed by alias name.

DES Event Services properties

Property	Required	Description
Alias	yes	Enter a unique name for this event source.
Start event source automatically on server startup		This option is set by default, which automatically starts this event source when MashZone NextGen Server starts. Clear this option if you need to manually control startup for this event source.
Event type	yes	Select the type of the event this event source should subscribe to. The XML schema files for these event types must exist in the Event Type Store directory of MashZone NextGen Event Service. The Event Type Store directory can be configured using Command Central.
Strategy	yes	The strategy that this event source uses for saving and removing events published from the Event Bus. Valid strategies are: <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until event source memory reaches capacity and then the event source removes the oldest events.

		<ul style="list-style-type: none"> ▪ DELTA= Events are stored by ID and added, updated or removed based on a command within the event. An event with an Insert command is saved in event source memory, any existing event with the same ID is overwritten. An event with a Remove command removes an existing event with the same ID.
Consider dimension		<p>Available only when Strategy is set to BUFFER. Set this option to save events in separate series (or buckets) for each unique value of the Dimension attribute.</p>
Dimension attribute	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Select the event attribute whose unique values determine separate event series (buckets) for this event source.</p>
Max. number of dimension values	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of unique dimension values (buckets) that can be tracked. Thus this is the maximum number of series that can store events.</p> <p>Default value: 1 (Max: 100.000).</p> <p>The product of Max. number of dimension values and Capacity per dimension value must not be greater than 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Determines how additional events are handled if they have new unique values for the dimension that defines buckets in this event source but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached.</p> <p>Set this option to change the bucket strategy to FIFO (first-in, first-out) which discards events for older series (buckets) and stores the newer event in a new</p>

		<p>series (bucket).</p> <p>Default value: false</p>
Capacity per dimension value	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value.</p> <p>Default value: 10</p> <p>The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.</p>
Event ID attribute	yes	<p>Available only when Strategy is set to DELTA.</p> <p>Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.</p>
Command attribute	yes	<p>Available only when Strategy if set to DELTA.</p> <p>Select the attribute that contains the event command (Insert or Remove). The event ID and command determines which events are stored, updated or removed in this event source.</p>
Capacity	yes	<p>Enter the maximum number of events to store in this event source. (Max: 100.000)</p> <p>Default value: 10</p>
Memory model		<p>Determines where events are stored:</p> <ul style="list-style-type: none"> ▪ Internal: the default which stores events in local memory for this event source. ▪ BigMemory: stores events in a local BigMemory cache.
Throttling		<p>Controls the speed and volume of event data that is pushed to views that subscribe to this event source. By default, event sources push event data every 500 milliseconds. You can:</p> <ul style="list-style-type: none"> ▪ Change the number of milliseconds to control throttling. ▪ Change the measurement (Default=500) to Events to have throttling wait until a specific number of events are received and change the number, if needed.

		See Example below.
Exception		<p>Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1):</p> <ul style="list-style-type: none"> ▪ A number ▪ Milliseconds or Events as the measurement for the exception criteria <p>See Example below.</p>

SIMPLE AND HYBRID THROTTLING STRATEGIES

Simple throttling strategies cause an event source to wait for either a specific time interval or for the receipt of a specific number of events and then push all new events to any subscribing real-time views. Throttling can slow event updates to real-time views when the volume or frequency for events causes rendering issues.

The default behavior is to push events to views based on a time interval of every 500 milliseconds. You can change the time interval or change the criteria to push events once a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies may still not even out event flow adequately. Instead, you can create hybrid strategies, such as "generally push every 50 milliseconds, but at most 10 events."

Hybrid strategies define the general throttling with the **Throttling** fields. You set the **Exception** option and define the exception that should break the general rule in the **Exception** criteria fields:

Throttling * Exception...

With the example hybrid throttling strategy shown above:

- The event source would wait 50 milliseconds after pushing events to subscribing views.

- If less than 10 events are received in that 50 milliseconds, they are pushed at the end of the interval.
- If a tenth event is received within the 50 milliseconds, these 10 events are pushed to subscribing views and both the time interval and the count of events begins again.
- If no events are received within the time interval, the event source waits until it receives an event. When an event is received, the event source pushes this event to subscribing views and restarts the time interval.

On the EDA Event Source overview page you can click on the **Alias** to show a preview of the specific Event Source properties.

8.6.5.2 Edit DES Event Sources

You can edit already existing DES Event Sources.

Note: Changes in DES connection properties can immediately affect data feed calculations so that they may not execute properly.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page is displayed.
4. Open the **Digital Event Services** tab.
5. Click the  **Edit** icon to configure a specific DES connection.
6. Set the properties for this event source. See the table below.
7. Click **Save**.

Your changes are applied.

Table 1. DES Event Services properties

Property	Required	Description
Alias	yes	Enter a unique name for this event source.
Start event source automatically on server startup		This option is set by default, which automatically starts this event source when MashZone NextGen Server starts. Clear this option if you need to manually control startup for this event source.
Event type	yes	Select the type of the event this event source should subscribe to. The XML schema files for these event types must exist in the Event Type Store directory of MashZone NextGen Event Service. The Event Type Store directory can be configured using Command Central.
Strategy	yes	The strategy that this event source uses for saving and

		<p>removing events published from the Event Bus. Valid strategies are:</p> <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until event source memory reaches capacity and then the event source removes the oldest events. ▪ DELTA= Events are stored by ID and added, updated or removed based on a command within the event. An event with an Insert command is saved in event source memory, any existing event with the same ID is overwritten. An event with a Remove command removes an existing event with the same ID.
Consider dimension		<p>Available only when Strategy is set to BUFFER.</p> <p>Set this option to save events in separate series (or buckets) for each unique value of the Dimension attribute.</p>
Dimension attribute	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Select the event attribute whose unique values determine separate event series (buckets) for this event source.</p>
Max. number of dimension values	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of unique dimension values (buckets) that can be tracked. Thus this is the maximum number of series that can store events.</p> <p>Default value: 1 (Max: 100.000).</p> <p>The product of Max. number of dimension values and Capacity per dimension value must not be greater than 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Determines how additional events are handled if they have new unique values for the dimension that defines buckets in this event source but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached.</p> <p>Set this option to change the bucket strategy to FIFO (first-in, first-out) which discards events for older series (buckets) and stores the newer event in a new series (bucket).</p> <p>Default value: false</p>

Capacity per dimension value	conditional	<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set.</p> <p>Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value.</p> <p>Default value: 10</p> <p>The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.</p>
Event ID attribute	yes	<p>Available only when Strategy is set to DELTA.</p> <p>Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.</p>
Command attribute	yes	<p>Available only when Strategy if set to DELTA.</p> <p>Select the attribute that contains the event command (Insert or Remove). The event ID and command determines which events are stored, updated or removed in this event source.</p>
Capacity	yes	<p>Enter the maximum number of events to store in this event source. (Max: 100.000)</p> <p>Default value: 10</p>
Memory model		<p>Determines where events are stored:</p> <ul style="list-style-type: none"> ▪ Internal: the default which stores events in local memory for this event source. ▪ BigMemory: stores events in a local BigMemory cache.
Throttling		<p>Controls the speed and volume of event data that is pushed to views that subscribe to this event source. By default, event sources push event data every 500 milliseconds. You can:</p> <ul style="list-style-type: none"> ▪ Change the number of milliseconds to control throttling. ▪ Change the measurement (Default=500) to Events to have throttling wait until a specific number of events are received and change the number, if needed. <p>See Example below.</p>
Exception		<p>Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1):</p> <ul style="list-style-type: none"> ▪ A number ▪ Milliseconds or Events as the measurement for the exception criteria

		See Example below.
--	--	--------------------

8.6.5.2.1 Simple and Hybrid Throttling Strategies

Simple throttling strategies cause an event source to wait for either a specific time interval or for the receipt of a specific number of events and then push all new events to any subscribing real-time views. Throttling can slow event updates to real-time views when the volume or frequency for events causes rendering issues.

The default behavior is to push events to views based on a time interval of every 500 milliseconds. You can change the time interval or change the criteria to push events once a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies may still not even out event flow adequately. Instead, you can create hybrid strategies, such as "generally push every 50 milliseconds, but at most 10 events."

Hybrid strategies define the general throttling with the **Throttling** fields. You set the **Exception** option and define the exception that should break the general rule in the **Exception** criteria fields:

Throttling * Exception...

With the example hybrid throttling strategy shown above:

- The event source would wait 50 milliseconds after pushing events to subscribing views.
- If less than 10 events are received in that 50 milliseconds, they are pushed at the end of the interval.
- If a tenth event is received within the 50 milliseconds, these 10 events are pushed to subscribing views and both the time interval and the count of events begins again.
- If no events are received within the time interval, the event source waits until it receives an event. When an event is received, the event source pushes this event to subscribing views and restarts the time interval.

8.6.5.3 Duplicate DES Event Sources

You can duplicate existing DES Event Sources.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page will be displayed.
4. Open the **Digital Event Services** tab.
5. Click the  **Copy** icon to duplicate a specific DES Event Source.

The selected DES Event Source is duplicated and listed with the prefix **copy_** in the **Alias**.

8.6.5.4 Delete DES Event Sources

You can delete existing DES Event Sources.

Warning

Deleting DES Event Sources may cause data feeds to fail.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page will be displayed.
4. Open the **Digital Event Services** tab.
5. Click the  **Delete** icon to delete a specific DES Event Source.

The selected DES Event Source is deleted from the list.

8.6.5.5 Share DES Event Sources

You can share DES Event Sources with particular users and user groups so that these have access to **DES Event Services**.

Prerequisites

You have administration privileges.

Regardless of the share, users with administration privilege can access all DES Event sources.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Digital Event Services** tab.

5. Click the  **Edit event service permissions** icon of the DES Event Service you want to share.
6. Enter a term in the search field and click **Search**. Clicking **Search** without any input values fetches all users and groups.
7. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search** results field.
8. Drag an user or user group from the Search result field and drop it into the **Principals with permissions** field.

By default, the owner of the **DES Event Services** is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.

9. Activate or deactivate the **Display** or **Usage** privileges of a user or user group.
A user or user group with **Display** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Usage** privilege has access to the relevant alias in the data source operator.
10. Click **OK**.

Your changes are applied.

8.6.5.6 Activate DES in MashZone NextGen

To use Digital Event Services (DES), a valid DES license file must be present in the MashZone NextGen installation.

The default path to the license is **<MashZone NextGen installation>/common/DigitalEventServices/license/license.xml**. After installation, a 30-days trial license is present in this location. To use DES, you must replace it by a valid license after 30 days.

8.6.6 Start or Stop an Event Source

To begin receiving events from the Event Bus, you must start the event source configured for that event type. You can also stop individual event sources.

Stopping an event source causes any existing events currently stored in memory to be deleted.

Procedure

Open the MashZone NextGen Admin Console (page 95).

1. Click **Event Service** to expand this section of the **Administration** menu.
2. Click **Event Service**. The **Event Service** page will be displayed.
3. Open the event source tab **EDA**, **Digital Event Services**, or **Apama** and either:
 - Select a specific event source and click  **Start** to start just that event source.
 - Or select a specific event source and click  **Stop** to stop that event source.

The selected event sources are stopped respectively started.

8.6.7 Restart all Event Sources

To begin receiving events from the Event Bus, you must start the event source configured for that event type. You can restart all event sources at once.

Note: Restarting all event sources causes any existing events currently stored in memory to be deleted.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**. The **Event Service** page will be displayed.
4. Click **Restart all** to restart all event source instances of the selected event source type.

All event sources are restarted.

8.6.8 Manage Apama Instances

By creating an Apama instance (Apama correlator) you can specify the connection to an Apama system.

You can create, edit and delete instances.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Apama instances**.
4. Select further steps:
 - Create Apama Instances (Seite 255)
 - Edit Apama Instances (Seite 256)
 - Delete Apama Instances (Seite 256)

8.6.8.1 Create Apama Instances

Creating an Apama instance (Apama correlator) you can specify the connection to an running Apama system.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Apama instances**.

4. Click **Create**.
5. Set the properties of the Apama instance. See table below.
6. Click **Save**.

The Apama instance is created and listed by alias name.

Apama instance properties

Property	Required	Description
Alias	yes	Enter a unique name for this Apama instance.
Host	yes	Host name to the running Apama system (local or remote)
Port	yes	Port number of the running Apama system (local or remote)

8.6.8.2 Edit Apama Instances

You can edit an already existing Apama instances.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Apama instances**.
4. Click the  **Edit** icon to configure an Apama instance.
5. Set the properties of the Apama instance. See table below.
6. Click **Save**.

Your changes are applied.

Apama instance properties

Property	Required	Description
Alias	yes	Enter a unique name for this Apama instance.
Host	yes	Host name to the running Apama system (local or remote)
Port	yes	Port number of the running Apama system (local or remote)

8.6.8.3 Delete Apama Instances

You can delete Apama instances.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Apama instances**.
4. Click the  **Delete** icon to delete a specific Apama instance.

The selected Apama instance is deleted from the list.

8.6.9 Manage Apama Event Targets

An Apama event target specifies an Apama system that can receive events sent by MashZone NextGen.

You can create, edit and delete Apama event targets.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Select further steps:
 - Create Apama Event Targets (Seite 257)
 - Edit Apama Event Targets (Seite 258)
 - Delete Apama Event Targets (Seite 259)
 - Share Apama Event Target (Seite 259)

8.6.9.1 Create Apama Event Targets

Creating an **Apama** event target you can specify an **Apama** system as target, receiving events from MashZone NextGen.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click **Create Apama Event Target**.
6. Set the properties for this event target. See table below.
7. Click **Save**.

The Apama event target is created and listed by alias name.

Property	Required	Description
Alias	yes	Enter a unique name for this event target.
Apama instance	yes	Alias with the pre-configured connection specification of a running Apama system (local or remote). See Manage Apama Instances (Seite 255) for details.
Event type	yes	Click  Refresh to update the list of Apama event types. Only event types are available that are present in the Apama instance selected. Select the type of the event this event target should subscribe to. Event types including not supported data types by MashZone NextGen are also available; these event types can be used for sending events, but the offending fields are not usable for data assignment to widgets.

8.6.9.2 Edit Apama Event Targets

You can edit an already existing Apama event target.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Edit** icon to configure an Apama event target.
6. Set the properties for this event target. See table below.
7. Click **Save**.

Your changes are applied.

Apama event target properties

Property	Required	Description
Alias	yes	Enter a unique name for this event target.
Apama instance	yes	Alias with the pre-configured connection specification of a running Apama system (local or remote). See Manage Apama Instances (Seite 255) for details.
Event type	yes	Click  Refresh to update the list of Apama event types. Select the type of the event this event target should subscribe to.

		Event types including not supported data types by MashZone NextGen are also available; these event types can be used for sending events, but the offending fields are not usable for data assignment to widgets.
--	--	--

8.6.9.3 Delete Apama Event Targets

You can delete Apama event targets.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Delete** icon to delete a specific Apama event target.

The selected Apama event target is deleted from the list.

8.6.9.4 Share Apama Event Target

You can share Apama event targets with particular users and user groups so that these have access to Apama event targets.

Prerequisite

You have administration privileges.

Regardless of the share, users with administration privilege can access all Apama event targets.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Event Service**.
4. Open the **Apama** tab.
5. Click the  **Edit permissions** icon of the Apama event target you want to share.
6. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
7. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
8. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

By default, the owner of the **Apama Event Services** is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.

9. Activate or deactivate the Display or Usage privileges of a user or user group.

A user or user group with **View** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Edit** privilege has access to the relevant alias in the data source operator.

10. Click **OK**.

Your changes are applied.

8.7 Process Performance Manager Integration

ARIS Process Performance Manager (PPM) lets you discover and analyze processes that are not formally managed by a business process management solution (**BPMS**), such as webMethods BPMS. Using data sources throughout your enterprise, such as transactional data from your business systems, event streams from webMethods BPMS or database records from trading partners, PPM can model a process and assess its performance across various dimensions, such as region, product line, volume, or time. You can also use PPM analytic tools to mine other data in your enterprise for meaningful patterns, trends, or correlations.

Information from PPM can be used as a source of data for MashZone NextGen dashboards and data feeds.

MashZone NextGen is compatible with PPM version 10.0 or above.

8.7.1 Manage PPM Connections

You can manage your PPM Connections in the **Admin console**.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **PPM connections** to expand this section of the **Administration** menu.
3. Click **PPM connections**.
4. Follow the procedure of the remaining steps:
 - Create PPM Connections (page 261)
 - Edit PPM Connections (Seite 262)
 - Delete PPM Connections (Seite 263)
 - Share PPM connections (Seite 263)

8.7.2 Create PPM Connections

You define connections for one or more PPM clients to allow users to use PPM as a data source for MashZone feeds or to allow users to add charts from PPM to workspace apps in MashZone NextGen.

MashZone NextGen is compatible with PPM 10.0 or above.

For MashZone NextGen to connect and retrieve PPM data or charts, the following PPM applications must be started:

- PPM
- PPM client

For details on your PPM installation, contact the system administrator in charge. You can enter PPM connection information manually or you can have MashZone NextGen determine them using the URL of a PPM favorite (favorites path). For information on copying the URL of a PPM favorite, see PPM on-line help topics.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **PPM connections** to expand this section of the **Administration** menu.
3. Click **PPM connections**.
4. Click **Create**.
5. Enter a name for the PPM connections in the Alias field, for example, the client name. The connection data is saved under this alias. Users may choose PPM connections by their alias.
6. To retrieve the connection data from the URL of a favorite from PPM:
 - a. Click **Retrieve data**.
 - b. Enter the URL of the PPM favorite that you copied earlier in the URL field.
 - c. Click **Resolve URL** to retrieve the required parameters from the URL. MashZone NextGen uses the favorite URL to complete the remaining fields for this connection.
7. To enter connection information manually:
 - a. Select the protocol (HTTP or HTTPS) to use for the web application server that hosts the PPM query interface.
For safety reason, we recommend to use the HTTPS protocol.
 - b. In the **Host** field, enter the fully qualified domain name of the PPM load balancer.
 - c. In the **Port** field, enter the port number of the PPM load balancer.
 - d. Specify the PPM client name of your PPM connection in the **Client** field.
8. Click **Check availability** to verify that the data is correct and that the PPM client is available.
9. Click **Save**.

The PPM connection is created and listed by alias name. This also lists the PPM version and availability of the PPM client.

8.7.3 Edit PPM Connections

You can edit already existing PPM connections.

Changes in PPM connection properties can immediately affect data feed calculations so that they may not execute properly.

For MashZone NextGen to connect and retrieve PPM data or charts, the following PPM applications must be started:

- PPM
- PPM client

For details on your PPM installation, contact the system administrator in charge. You can enter PPM connection information manually or you can have MashZone NextGen determine them using the URL of a PPM favorite (favorites path). For information on copying the URL of a PPM favorite, see PPM on-line help topics.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **PPM connections** to expand this section of the **Administration** menu.
3. Click **PPM connections**. A list of all available PPM connections is displayed.
4. Click the  **Edit** icon to configure a PPM connection.
5. The **Alias** field of an already configured PPM connection is not editable. The connection data is saved under this alias.
6. To retrieve the connection data from the URL of a favorite from PPM:
 - a. Click **Retrieve data**.
 - b. Enter the URL of the PPM favorite that you copied earlier in the **URL** field.
 - c. Click **Resolve URL** to retrieve the required parameters from the URL. MashZone NextGen uses the favorite URL to complete the remaining fields for this connection.
7. Enter the connection information manually.
 - a. Select the protocol (HTTP or HTTPS) to use for the web application server that hosts the PPM query interface.
For safety reason, we recommend to use the HTTPS protocol.
 - b. In the **Host** field, enter the fully qualified domain name of the PPM load balancer.
 - c. In the **Port** field, enter the port number of the PPM load balancer.
 - d. Specify the PPM client name of your PPM connection in the **Client** field.
8. Click **Check availability** to verify that the data is correct and that the PPM client is available.
9. Click **Save**.

Your changes are applied.

8.7.4 Delete PPM Connections

You can delete existing PPM connections.

Warning

Deleting PPM connections may cause data feeds to fail.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **PPM connections** to expand this section of the **Administration** menu.
3. Click **PPM connections**. A list of all available PPM connections is displayed.
4. Click the  **Delete** icon to delete a PPM connection.
5. Confirm the deletion.

The selected PPM connections are deleted from the list.

8.7.5 Share PPM connections

You can share PPM connections with particular users and user groups so that these have access to PPM server.

You have administration privileges.

Regardless of the share, users with administration privilege can access all PPM connections.

Regardless of the share, users with administration privilege can access all EDA Event sources.

Procedure

1. Open the MashZone NextGen Admin Console (page 95).
2. Click **PPM Connections**.
3. Click the  **Edit PPM alias permissions** icon of the PPM connection you want to share.
4. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
5. Click **Show MashZone NextGen default groups** to show only default MashZone NextGen users or user group in the **Search results** field.
6. Drag an user or user group from the **Search result** field and drop it into the **Principals with permissions** field.

By default, the owner of the PPM connection is already present in the **Principals with permissions** list . This owner is non editable and cannot be removed from the list.

7. Activate or deactivate the **Display** or **Usage** privileges of a user or user group.

A user or user group with **Display** privilege can see the relevant source data in the data feed or dashboard. A user or user group with the **Usage** privilege has access to the relevant alias in the data source operator.

8. Click OK.

Your changes are applied.

8.8 webMethods Business Console Integration

MashZone NextGen can be easily embedded in webMethods Business Console using a native Business Console gadget.

The gadget is called MashZone NextGen and can be found in the Business Console **Common** section.

Detailed information on how to use webMethods Business Console can be found in the documentation **Working with webMethods Business Console**.

In order to access a MashZone NextGen dashboard, the **Dashboard URL** must be provided in the gadget settings. The URL must contain the MashZone NextGen dashboard GUID as an URL parameter.

8.8.1 Example

`http://sbrvpresto4.eur.ad.sag:8080/mashzone/hub/dashboard/dashboard.jsp?guid=e35c1619-0b06-42ac-b343-b16e7d5dcc12`

In the section **UI Settings** the gadget height, title and border can be specified.

The section **Data Mapping** specifies the parameters required for the communication between MashZone NextGen dashboards and Business Console gadgets.

- **Mapping Id:** Identifier used in Business Console for gadget to gadget communication. The **Mapping Id** is needed to identify and map the data send from one gadget to the data structure of another gadget. In case of two MashZone NextGen gadgets, exchanging data, the mapping can be used to take a selection value from one embedded MashZone NextGen widget and use it as selection value in the other MashZone NextGen widget.
- **Widget Id:** Specifies the external identifier of the MashZone NextGen widget to communicate with.
- **Widget Parameter:** Specifies a measure or dimension name used in the MashZone NextGen widget.
- **Default Value:** Optionally, it is possible to define a default value, that is used for example as a default selection for the MashZone NextGen widget after loading the gadget in Business Console.

All data required can be found in MashZone NextGen, see **Use dynamic URL selection** for details.

8.8.2 Authentication

You can integrate MashZone NextGen under My webMethods in an SSO scenario by SAML (Security Assertion Markup Language).

MashZone NextGen can accept SAML tokens for authentication in a SSO environment.

See Authentication with Single Sign-On Solutions (page 100) for details.

A BASE64 encoded SAML token is expected. Since it is send via URL it needs to be URL encoded before.

8.8.3 Example URL

```
http://sbrvpresto4.eur.ad.sag:8080/mashzone/hub/dashboard/dashboard.jsp?appheader=fals
e&guid=64545b4f-d150-4241-a858-2304eea23684 &SAMLToken=<URL encodedBASE64
encoded token>
```

8.8.4 Configuration

To enable access to MashZone NextGen you need to list the URL(s) of the webMethods Business Console server(s) in the Content Security Policy of MashZone NextGen.

The content security settings are done in the server configuration file

applicationContext-security-filters.xml by adding filters for X-Frame-Options and Content Security Policies. The file is located in <MashZone NextGen installation>\apache-tomcat\webapps\mashzone\WEB-INF\classes.

applicationContext-security-filters.xml (abstract)

```
<beans:beans xmlns="http://www.springframework.org/schema/security"...> ... <http
pattern="/hub/login.html" security="none"/> <http pattern="/**/*.*.jsp"
use-expressions="false" authentication-manager-ref="authenticationManager"
entry-point-ref="mzngAuthenticationEntryPoint"> <anonymous enabled="false"/>
<headers> <!--frame-options policy="SAMEORIGIN"/--> <frame-options
policy="ALLOW-FROM" strategy="whitelist"
value="http://BCServerHostA:BCServerPortA,
http://BCServerHostB:BCServerPortB,..."/> <!--content-security-policy
policy-directives="frame-ancestors 'self'"/--> <content-security-policy
policy-directives="frame-ancestors 'self' http://BCServerHostA:BCServerPortA,
http://BCServerHostB:BCServerPortB,.../> </headers> <csrf
token-repository-ref="csrfTokenRepository"
request-matcher-ref="skipHttpAuthCsrfMatcher"/> </http> <http pattern="/**/*.*.html"
use-expressions="false" authentication-manager-ref="authenticationManager"
entry-point-ref="mzngAuthenticationEntryPoint"> <intercept-url
pattern="/**/*.*.html" access="IS_AUTHENTICATED_ANONYMOUSLY"/> <anonymous
enabled="false"/> <headers> <!--frame-options policy="SAMEORIGIN"/-->
<frame-options policy="ALLOW-FROM" strategy="whitelist"
value="http://BCServerHostA:BCServerPortA,
http://BCServerHostB:BCServerPortB,..."/> <!--content-security-policy
policy-directives="frame-ancestors 'self'"/--> <content-security-policy
policy-directives="frame-ancestors 'self' http://BCServerHostA:BCServerPortA,
http://BCServerHostB:BCServerPortB,..."/> </headers> </http> ... </beans:beans>
```

8.8.5 Outbound API

MashZone NextGen provides an outbound API to pass data from MashZone NextGen dashboards to an embedding system, for example, an external web application like webMethods Business Console.

See **Post data** (page 66) for details.

8.8.6 Inbound API

By using iFrame MashZone NextGen can be used as a widget in external products, for example, webMethods Business Console. As embedded widget MashZone NextGen is enabled to send data via outbound API (Post data) to the embedding system and receive data via inbound API (URL selection) from the embedding system.

See Embedding MashZone NextGen in external system environments (page 155) for details.

8.9 MashZone NextGen Repositories

The MashZone NextGen Repository is the database that the MashZone NextGen Server uses to store meta-data, attributes and configuration for MashZone NextGen including:

- Artifacts (dashboards and data feeds)
- Configuration properties for the MashZone NextGen Server

If you are using the default User Repository, user and group data is also stored in the MashZone NextGen Repository.

The MashZone NextGen repository is initially installed in a Derby database suitable only for trial purposes. For proof-of-concept, development or production uses, move the repositories to a robust and compatible solution.

Configuration and administration tasks for these two repositories include:

- Move the MashZone NextGen repository to a robust database solution (page 116)
- Support International Character Sets and Locales (page 150)
- Use the Default MashZone NextGen User Repository (page 132)
- Change MashZone NextGen Repository Ports (page 154)
- Tuning the MashZone NextGen Repository Connection Pool (page 267)
- Synchronize the MashZone NextGen Repository and MashZone NextGen Server Time Zones (page 268)
- Sharing the MashZone NextGen Repository in Clustered Environments (page 212)
- Configure BigMemory Servers for MashZone NextGen Caching and In-Memory Stores (page 170)
- Maintenance Suggestions (page 266)

8.9.1 Maintenance Suggestions

Your existing standards for database backups, security and maintenance can be applied to the MashZone NextGen repositories. In addition, you should set up procedures to monitor or

regularly manage growth for the MashZone NextGen Auditable Events table. This table tracks audit information for updates to the MashZone NextGen Repository.

You may also want to move snapshot data to a separate database to more easily manage growth and other operations for these datasets.

8.9.2 Tuning the MashZone NextGen Repository Connection Pool

In addition to basic connection configuration, you can configure the connection pools for the MashZone NextGen Repository. In many cases, you need to tune this configuration to optimize your MashZone NextGen environments.

For a complete list of connection properties, see Tomcat Datasource Properties (<http://commons.apache.org/proper/commons-dbcp/configuration.html>).

To tune the connection pool, you update properties in the **<Resource>** element for the MashZone NextGen repository in the `<MashZone NextGen installation>/apache-tomcat/conf/context.xml` file and then restart MashZone NextGen to apply these changes.

8.9.2.1 Connection Pool Size Properties

initialSize	The initial number of connections to create when the pool starts up. This defaults to 0.
maxWaitMillis	The maximum number of milliseconds that the pool will wait when no connections are available before failing. Defaults to -1 which is an indefinite wait.

8.9.2.2 Idle Pool Connection Properties

maxIdle	The maximum number of connections that can be idle without connections being released. Defaults to 20. Set this to -1 to prevent any connections being released.
minIdle	The minimum number of idle connections that can exist before new connections are added to the pool. This defaults to 0, indicating no new connections should be created.
testWhileIdle	Whether connections should be tested when idle. If this is enabled, idle connections are tested using the Validation query . See Move the MashZone NextGen repository to a robust database solution (page 116) for more information on

	validation queries.
timeBetweenEvictionRunsMilliseconds	The number of milliseconds between tests of idle connections. This defaults to -1, which prevents all idle connection testing.
numTestsPerEvictionRun	The number of connections to test during any idle connection test run.
minEvictableIdleTimeMillis	The minimum number of milliseconds that a connection can be idle before being tested for eviction. Default is 3 minutes.

8.9.3 Synchronize the MashZone NextGen Repository and MashZone NextGen Server Time Zones

Creation and modification timestamps for artifacts and other MashZone NextGen Repository metadata can be different than times when events occurred in the MashZone NextGen Server in two cases:

- If the server hosting the MashZone NextGen Repository is located in a different time zone from the server hosting the MashZone NextGen Server
- If the time zone setting for the database hosting the MashZone NextGen Repository is set to a different time zone from the server hosting the MashZone NextGen Server

You can correct this problem by specifying a time zone in configuration for the MashZone NextGen Repository.

The instructions in this topic are specific to MySQL databases. For other types of databases, please consult documentation for that database to determine the appropriate updates.

9 Appendix

9.1 Widgets

The dashboard editor provides a number of widgets.

See also [Insert widgets in a dashboard \(page 30\)](#).

9.1.1 Action button

Can be used to trigger an action.

The following widget options are available.

General options	Description
Name	Optional widget name.
Container	<p>Hide header: Hides the header as well as the title of the component, and resizes the content of the container. Click the icon again to display the header. The header is hidden by default.</p> <p>Hide border: Hides the outline of the component container. Click the icon again to display the outline. The border is visible by default.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Style	Selects the style type of the component. The component styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.

Specific bar chart options	Description
Label	Label that is displayed in the widget.

9.1.2 Bar chart

A bar chart can display values for two iterations:

- Two dimensions and one KPI
- One dimension and multiple KPIs

The second iteration is displayed in the form of several stacked bars. If multiple KPIs are used, these are displayed in bars of different colors.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific bar chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
Column type	Selects the bar type of the bar chart. This option is only available if a data source has been assigned. Bars available are Stacked or Grouped .
Legend position	Displays a legend in the widget and sets the legend's position. The default is None , that is, no legend is displayed.
Axis	<p>Specifies the visualization of the X- and Y-axis.</p> <p>Show all: Displays the axis titles and labels. Enabled by default.</p> <p>Hide title: Hides the axis titles and displays only the axis labels.</p> <p>Hide all: Hides the axes completely.</p>

Specific bar chart options	Description
Multiple selection	In view mode, the user can select multiple values at the same time in the corresponding widget, for example, multiple rows in a table or multiple coordinates in a chart. The multiple selection can be used, for example, to filter other dashboard widgets. See Use interactive filters in dashboards (Seite 5) for details. The values selected are processed as a list. If the option is enabled, columns of List type are provided in the filter configuration dialog to configure filter conditions. See Define filters for dashboard widgets (Seite 60) for details.

9.1.3 Bubble chart

A bubble chart displays one dimension and two KPIs. The two KPIs are plotted on the X- and Y-axis. The dimension is represented by different colors of the individual bubble areas. Optionally, a third KPI can be incorporated; its values determine the radii of the bubble areas.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific bubble chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
Data points	Selects the data point shape. The default is Circle . The option is available only if a Partition is assigned. See Assign data columns to bubble chart. (page 37)
Data point size	Selects the data point size. The default is Medium .
Legend position	Displays a legend in the widget and sets the legend's position. The default is None , that is, no legend is displayed.
Axis	Specifies the visualization of the X- and Y-axis. Show all : Displays the axis titles and labels. Enabled by default. Hide title : Hides the axis titles and displays only the axis labels. Hide all : Hides the axes completely.

9.1.4 Circular gauge chart

A circular gauge chart displays a set of aggregated KPI values. The value ranges are arranged in a semicircle with a red pointer being displayed that indicates the actual value of the KPI.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	Hide header : Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header. Hide border : Hides the outline of the widget container. Click the icon again to display the outline.
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.

General options	Description
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific circular gauge chart options	Description
Scale	Displays a scale in the chart.
Scale value	Displays the scale values in the chart.
KPI name	Displays the KPI name in the chart.
KPI value	Displays the KPI value in the chart.
Threshold	Displays the threshold in the chart.
Threshold value	Displays the threshold value in the chart.
Level meter	Displays a level meter in the chart.
Percentage	Displays a scale from 0 to 100% in the chart. The KPI value is also displayed in percent.

9.1.5 Column chart

A column chart can display values for two iterations:

- Two dimensions and one KPI
- One dimension and multiple KPIs

The second iteration is displayed in the form of several stacked columns. If multiple KPIs are used, these are displayed in columns of different colors.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header. Hide border: Hides the outline of the widget container. Click the icon

General options	Description
	again to display the outline.
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific column chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
Rotate X-axis labels	Rotates the labels of the X-axis 90 degrees.
Column type	Selects the column type of the column chart. This option is only available if a data source has been assigned. Columns available are Stacked or Grouped .
Legend position	Displays a legend in the widget and sets the legend's position. The default is None , that is, no legend is displayed.
Axis	Specifies the visualization of the X- and Y-axis. Show all : Displays the axis titles and labels. Enabled by default. Hide title : Hides the axis titles and displays only the axis labels. Hide all : Hides the axes completely.
Multiple selection	In view mode, the user can select multiple values at the same time in the corresponding widget, for example, multiple rows in a table or multiple coordinates in a chart. The multiple selection can be used, for example, to filter other dashboard widgets. See Use interactive filters in dashboards (Seite 5) for details. The values selected are processed as a list. If the option is enabled, columns of List type are provided in the filter configuration dialog to configure filter conditions. See Define filters for dashboard widgets (Seite 60) for details.

9.1.6 Date filter

Provides an interactive calendar for configuring a date filter. The calendar filters the values in a given date range.

In view mode, you can adjust the date range by clicking the + symbol and - symbol in the widget.

The **Date filter** widget supports values only of date type.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific date filter options	Description
Granularity	<p>Specifies the structure of the time range.</p> <ul style="list-style-type: none"> ▪ There are three possible time granularities: ▪ Year ▪ Quarter ▪ Month
Allow range selection	<p>Enables the selection of multiple time periods in the calendar.</p> <p>To define a time range in view mode, select multiple periods by holding the mouse button.</p> <p>The option is enabled by default. If the option is disabled, the</p>

Specific date filter options	Description
	From and To values of the default selection are set to the same value.
Compact	Hides the selected range label above the time filter.
Range	<p>Specifies the initially displayed date range with a start value (From) and an end value (To).</p> <p>You can either enter the date value manually or extract it from a data feed. This option is disabled for manual entry if the date value is extracted from a data feed.</p> <p>If no data source columns are assigned to the widget, you can enter the initial range values manually.</p> <p>If data source columns are assigned, the range values are taken from the corresponding columns. The first values of the corresponding data source columns are always taken as the start or the end value of the range.</p>
Default selection	<p>Specifies the range values that are preselected by default.</p> <p>You can either enter the date value manually or extract it from a data feed. This option is disabled for manual entry if the date value is extracted from a data feed.</p> <p>If no data source columns are assigned to the widget, you can enter the default range values manually.</p> <p>If data source columns are assigned, the default values are taken from the corresponding columns.</p> <p>The first values of the corresponding data source columns are always taken as the minimum or the maximum.</p>

9.1.7 Drop-down box

The drop-down box provides you with a selection of values in a drop-down menu for you to filter other widgets.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.

General options	Description
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific drop-down box options	Description
No selection	Adds the No selection value to the value list as the first entry. By default, the first entry of the value list is preselected in the drop-down box for filtering. Selecting the No selection value has no effect on filtering other widgets.
No selection label	Alternative text for No selection default. The text is displayed in the drop-down box.

9.1.8 Grid

You can use the **Grid** widget to insert a table in your dashboard.

The following widget options are available.

General options	Description
Name	Optional component name.
More options	Displays additional options.

General options	Description
Container	<p>Hide header: Hides the header as well as the title of the component, and resizes the content of the container. Click the icon again to display the header. The header is hidden by default.</p> <p>Hide border: Hides the outline of the component container. Click the icon again to display the outline. The border is visible by default.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Style	Selects the style type of the component. The component styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the component.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific grid options	Description
Edit columns	Enables you to edit the Grid widget interactively. Click the Configuration icon. In edit mode, the Grid widget has a blue frame. You can set the initial widths and the initial sort order of the columns. You can change the sort order by clicking the column header. To adapt the column width, drag the column borders using the mouse. You can change the defined column width and column sort order in dashboard view mode (Seite 8).
Auto column width	Automatically adapts the column width to the column content. The horizontal scroll bar is no longer displayed in the widget.
Multiple selection	In view mode, the user can select multiple values at the same time in the corresponding widget, for example, multiple rows in a table or multiple coordinates in a chart. The multiple selection can be used, for example, to filter other dashboard widgets. See Use interactive filters in dashboards (Seite 5) for details. The values selected are processed as a list. If the option is enabled, columns of List type are provided in the filter configuration dialog to configure filter conditions. See Define filters for dashboard widgets (Seite 60) for details.

9.1.9 Horizontal and vertical gauge chart

A gauge chart displays a set of aggregated KPI values. The value ranges are arranged in a horizontal or vertical bar with a pointer being displayed that indicates the actual value of the KPI. The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific circular gauge chart options	Description
Scale	Displays a scale in the chart.
Scale value	Displays the scale values in the chart.
KPI name	Displays the KPI name in the chart.
KPI value	Displays the KPI value in the chart.
Threshold	Displays the threshold in the chart.
Threshold value	Displays the threshold value in the chart.
Level meter	Displays a level meter in the chart.
Percentage	Displays a scale from 0 to 100% in the chart. The KPI value is also displayed in percent.

Specific circular gauge chart options	Description
Scale size	Selects the scale size of the gauge chart. The default is Medium .

9.1.10 Image

The **Image** widget supports the file format PNG, GIF, and JPEG. The image file can be selected using a Web URL or a URL Alias.

For information about specifying the image source URL, see Specify an image source URL (page 65).

General options	Description
Name	Optional component name.
Container	<p>Hide header: Hides the header as well as the title of the component, and resizes the content of the container. Click the icon again to display the header. The header is hidden by default.</p> <p>Hide border: Hides the outline of the component container. Click the icon again to display the outline. The border is visible by default.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Style	Selects the style type of the component. The component styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.

Specific image options	Description
Image sizing	<p>Original: Displays the image in the original size.</p> <p>Scaled: Fits the display of the image to the widget frame size.</p> <p>Aspect ratio: Keeps the aspect ratio of the image.</p>
Horizontal alignment	Aligns the image horizontally within the widget frame.
Vertical alignment	Aligns the image vertically within the widget frame.

9.1.11 Input field

The input field enables you to manually enter values in order to filter other widgets.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific input field options	Description
Data type	<p>Data type of the user input. The user can insert only values of the selected data type.</p> <ul style="list-style-type: none">  Text  Number  Date <p>If the input field data type is set to DATE, a  calendar icon that you can use to open a date picker is shown in the input field.</p>
Date format	<p>Format of the user input. The user can insert only values that comply with the selected format.</p> <p>The option is available only for data type Date.</p>

Specific input field options	Description
Prompt text	Displays a prompt text in the input field.
Initial value	Initial value is displayed in the input field and used as preset value. Optional
Current date/time	Uses the current date as the value that is preset by default. The widget always uses the current date and time of the dashboard at runtime. The option is available only for data type Date . The value displayed depends on the format selected.
Submit button text	Alternative text for the submit button. The default text is Ok .
Show submit button	Displays the submit button in the input field.
Submit value with each keystroke	Immediately submits the entered values with each keystroke. The input field only submits valid date or number values for Date and Number .

9.1.12 Label

A label displays a fixed text you have entered, or a text that is supplied dynamically by a data source.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header. Hide border: Hides the outline of the widget container. Click the icon again to display the outline.
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.

General options	Description
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific label options	Description
Style	Style to be used for the text displayed in the drop-down menu.
Alignment	Alignment of the text that is displayed. The default is left aligned .

9.1.13 Line chart

A line chart can display values for two iterations:

- Two dimensions and one KPI
- One dimension and multiple KPIs

The second iteration is displayed in the form of several stacked lines. If multiple KPIs are used, these are displayed in lines of different colors.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header. Hide border: Hides the outline of the widget container. Click the icon again to display the outline.
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display

General options	Description
	the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific line chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
Line type	Selects the line type of the line chart. Linear , Curved or Step line types are available. The default is Linear .
Data points	Selects the size of displayed data points. Line types available are Large , Small or None . The default is Large .
Interpolate	Enabled: A lack of values results in a continuous line. Disabled: A lack of values results in gaps within the line.
Rotate X-axis labels	Rotates the labels of the X-axis 90 degrees.
Legend position	Displays a legend in the widget and sets the legend's position. The default is None , that is, no legend is displayed.
Axis	Specifies the visualization of the X- and Y-axis. Show all : Displays the axis titles and labels. Enabled by default. Hide title : Hides the axis titles and displays only the axis labels. Hide all : Hides the axes completely.

9.1.14 List

The **List** widget lists the values of one or two assigned data source columns.

You can use the **List** widget to select values, for example, to filter multiple values in other widgets.

In view mode, the widget provides check boxes for selection when multiple selection is enabled. You can select multiple values by pressing the **Shift** key and clicking several rows.

The following widget options are available.

General options	Description
Name	Optional component name.

General options	Description
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the component, and resizes the content of the container. Click the icon again to display the header. The header is hidden by default.</p> <p>Hide border: Hides the outline of the component container. Click the icon again to display the outline. The border is visible by default.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Style	Selects the style type of the component. The component styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the component.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific list options	Description
Edit columns	Enables you to edit the List widget interactively. Click the Configuration icon. In edit mode, the List widget has a blue frame. You can set the initial widths and the initial sort order of the columns. You can change the sort order by clicking the column header. To adapt the column width, drag the column borders using the mouse. You can change the defined column width and column sort order in dashboard view mode (Seite 8).
Multiple selection	<p>In view mode, the user can select multiple values at the same time in the corresponding widget, for example, multiple rows in a table or multiple coordinates in a chart. The multiple selection can be used, for example, to filter other dashboard widgets. See Use interactive filters in dashboards (Seite 5) for details. The values selected are processed as a list. If the option is enabled, columns of List type are provided in the filter configuration dialog to configure filter conditions. See Define filters for dashboard widgets (Seite 60) for details.</p> <p>The option is enabled by default.</p>

9.1.15 Map with markers

A map with markers allows you to mark points of interest on a map based on geo-coordinates (latitude & longitude) defined in the assigned data.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific map with markers options	Description
Base map	<p>Selects a base map. A base map provides you with context for a map. You can add information to a base map by overlaying other information on top of it.</p> <p>Default base map is Open Street Maps.</p> <p>Select a base map in the drop-down menu.</p> <p>You can add your own base maps. For details, see Manage tile server configuration files (page 187).</p>
Default marker icon	<p>Selects the default marker icons.</p> <p>Default icon is a circle.</p>

Specific map with markers options	Description
	Select an icon in the drop-down menu.
Sticky active area	Adjust automatically the zoom level of the map based on the active markers to display all markers in the widget.
Show label	Displays the label of all regions if applicable.

9.1.16 Pie chart

A pie chart can display one numerical KPI iterated over a dimension (text or date dimension). The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific pie chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
Inner radius	Specifies a inner radius in % to display an inner circle using the specified radius.
Values	Selects the display type of the KPIs.
Values position	Selects the position of the displayed KPI values. Positions available are Inside and Outside .
Legend position	Displays a legend in the widget and sets the legend's position. The default is None , that is, no legend is displayed.
Multiple selection	In view mode, the user can select multiple values at the same time in the corresponding widget, for example, multiple rows in a table or multiple coordinates in a chart. The multiple selection can be used, for example, to filter other dashboard widgets. See Use interactive filters in dashboards (Seite 5) for details. The values selected are processed as a list. If the option is enabled, columns of List type are provided in the filter configuration dialog to configure filter conditions. See Define filters for dashboard widgets (Seite 60) for details.

9.1.17 Rich text area

The **Rich text area** widget displays a fixed text that you have entered, or a text that is supplied dynamically by a data source or a selected element of another widget.

The widget provides a text editor in which you can enter and format text. You can also insert variable data fields to dynamically display values of a data source.

The **Rich text area** widget supports multi-lingual dashboards (page 27). You can translate the widget title into different languages, but not the contents of the widget.

Double-click the widget to open the text editor.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon</p>

General options	Description
	again to display the outline.
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific Rich text area options	Description
Text style	<p>Provides options for text layout.</p> <ul style="list-style-type: none"> ▪ Text format, for example, Heading, Titel, or Normal, provided in the drop-down menu. ▪ Increase and decrease text size ▪ Bold ▪ Italic ▪ Underline ▪ Strikethrough ▪ Text color ▪ Clear text style
Paragraph	<p>Provides options for paragraph layout.</p> <ul style="list-style-type: none"> ▪ Text alignment ▪ Text indent ▪ Ordered and unordered list
Insert dynamic values	<p>Inserts variable data fields at the cursor position. The variable data fields allow you to display values that are dynamically supplied by assigned columns of the data source. See Assign data columns to rich text editors.</p>

9.1.18 Slider

Provides the user with a selection of values in the form of a slider. The slider filters the values in a given data range. You can set the maximum and minimum of the data range using indicators. The selected minimum and maximum values of the slider are labeled. In view mode, you can manually edit the minimum and maximum values of the labels. This allows you to enter exact values for the selected data range. Click a label of an indicator and enter a value required. In addition, you can adjust the minimum and maximum values individually or you can move the entire selected data range using the mouse pointer.

The **Slider** widget supports only numeric values.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific slider options	Description
Type	<p>Selects one of the following slider types:</p> <ul style="list-style-type: none"> ▪ Single value: With this slider, you can select a single value in a defined range of values. ▪ Range: With this slider, you can select the minimum (From) and the

Specific slider options	Description
	<p>maximum (To) value of a data range.</p> <ul style="list-style-type: none"> ▪ Range (fixed upper): With this slider, you can select the minimum (From) value of a data range. The maximum value is fixed. ▪ Range (fixed lower): With this slider, you can select the maximum (To) value of a data range. The minimum value is fixed.
Display range	<p>Specifies the data range with a minimum value (From) and a maximum value (To).</p> <p>If no data source columns are assigned to the widget, you can enter the range values manually.</p> <p>If data source columns are assigned, the range values are taken from the corresponding columns. The first values of the corresponding data source columns are always taken as minimum or maximum.</p> <p>The first value of the column assigned to the data range maximum should be greater than the first value of the column assigned to the data range minimum.</p> <p>The preset values are 0 (minimum) and 100 (maximum).</p>
Initial selection	<p>Values that are preselected by default.</p> <p>For the Single type, you specify a single value.</p> <p>For the Range type, you specify the From and To values of the data range.</p> <p>For the Range (fixed upper) type, you specify the From value.</p> <p>For the Range (fixed lower) type, you specify the To value.</p> <p>If no data source columns are assigned to the widget, you can enter the initial range values manually.</p> <p>If data source columns are assigned, the initial values are taken from the corresponding columns. The first values of the corresponding data source columns are always taken as minimum or maximum.</p> <p>By default, the slider selects the full available range.</p>
Show ticks	<p>Shows the step ticks.</p> <p>Enabled by default.</p>

Specific slider options	Description
Step width	<p>Step width of the values between minimum and maximum value.</p> <p>The values of the slider, starting with the minimum, are increased by this value until the maximum is reached.</p> <p>By default, the step width is set automatically.</p> <p>The minimum permissible step width is not smaller than 1/1000th of the range value. For example, if the range is 0 to 10000, the minimum step value is 10.</p> <p>You can enter your own value for the step width.</p> <p>The option is only available if the Show ticks option is enabled.</p>
Only step values	<p>Allows the user to select step values only. The indicator snaps to the nearest tick (step value) if the indicator is left between two ticks.</p> <p>Disabled by default.</p> <p>The option is only available if the Show ticks option is enabled.</p>
Enable markers	<p>Displays the selected slider value in the marker.</p> <p>The option is enabled by default.</p>
Rotate axis	<p>Displays the values at an angle of 45° on the axis if the option is enabled, otherwise otherwise displays the values horizontally.</p> <p>The option is disabled by default.</p>
Numeric format	<p>Selects the format of the values displayed in the widget.</p> <p>The slider widget supports only numeric values.</p>

9.1.19 Jump to PPM client

The **Jump to PPM client** widget displays a customizable text which can be used to jump to PPM. Optionally, you can configured a favorite which is displayed when you jump to PPM.

The following widget options are available.

General options	Description
Name	Optional component name.

General options	Description
Container	<p>Hide header: Hides the header as well as the title of the component, and resizes the content of the container. Click the icon again to display the header. The header is hidden by default.</p> <p>Hide border: Hides the outline of the component container. Click the icon again to display the outline. The border is visible by default.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.
Style	Selects the style type of the component. The component styles available in the drop-down menu are part of the style template selected for the current dashboard. The Default widget style is preselected.

Specific Jump to PPM widget options	Description
PPM Connections	Alias of the PPM Connection, which contains the PPM client connection data defined in MashZone NextGen.
Favorite	Path of a PPM favorite. The favorite path represents the favorites tree including favorites folder and name, for example, \Favorites\Process cycle time.
Extract from URL	Automatically determines the connection data of the PPM data source, for example, alias, favorite path, and favorite type, click Extract from URL and insert the favorite URL created in PPM.
Favorite type	Favorite type specifies the favorite as Private or Shared .

SAML AUTHENTICATION

If a user starts PPM using the Jump to PPM widget, he is automatically authenticated against the PPM system using single sign-on (SSO), provided SSO is configured correctly for PPM and ARIS Connect. The current logon language of the user is used for PPM, if available. Otherwise, the default language of PPM is used.

There are several ways the PPM client can be open:

- PPM applet always opens in a new tab.
- PPM web start opens a new client window (If no favorite jump client window is open.).
- PPM web start reuses an open client window (If a favorite client window is open, and the languages match.).

- PPM web start opens a new windows and closes an open one (If a favorite client window is open, and the languages do not match.).

9.1.20 Traffic lights

A multi-color vertical, horizontal, or single traffic light shows the threshold range in which a KPI value is located. You can define thresholds in the **Assign data** dialog.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific traffic light options	Description
Type	Types available are Vertical , Horizontal and Single .

9.1.21 Vector map

A vector map allows you to interact with a vector-based map and to visualize geographic areas of interest.

The following widget options are available.

General options	Description
Name	Optional widget name.
More options	Displays additional options.
Container	<p>Hide header: Hides the header as well as the title of the widget , and resizes the content of the container. Click the icon again to display the header.</p> <p>Hide border: Hides the outline of the widget container. Click the icon again to display the outline.</p>
Container style	Selects the style type of the container. The container styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Style	Selects the style type of the widget . The widget styles available in the drop-down menu are part of the style template selected for the current dashboard. By default, the Default widget style is preselected.
Auto refresh	Sets the automatic data retrieval of the widget.
Show menu	Enables the widget menu in view mode. In view mode, you can display the widget menu by clicking the  Menu icon in the widget header. In the widget menu, you can, for example, clear a selection, refresh the data displayed, or save the component widget in a CSV file.

Specific vector map options	Description
Template	<p>Selects a map template of outline maps based on available GeoJSON configurations.</p> <p>Default map template is world-countries-by-name.</p> <p>Select a map template in the drop-down menu.</p> <p>You can add your own map templates. For details, see Manage geoJSON files (page 186).</p>
Base map	Selects a base map. A base map provides you with context for a map. You can add information to a base map by overlaying other information on top of it.

Specific vector map options	Description
	Default base map is Open Street Maps. Select a base map in the drop-down menu.
Sticky active area	Adjust automatically the zoom level of the map based on the active regions to display all regions in the widget.
Show label	Displays the label of all regions if applicable.

9.2 Operators

MashZone NextGen provides a wide range of data source and data transformation operators for creating dashboards and data feed definitions.

- Data source operators (page 296)
A data source operator enables you to specify the connection to a data source and to configure the data retrieval.
- Data transformation operators (page 317)
In addition to the data source operators you can add operators to transform the source data. The following data transformation operators are available in the MashZone NextGen.
- User input operators (page 357)
User input operators enables the dynamic entry of single values or value lists in the data feed processing.

9.2.1 Data source operators

A data source operator enables you to specify the connection to a data source and to configure the data retrieval.

The following data source operators are available in the dashboard and data feed editor.

9.2.1.1 ARIS table

Extracts data from an ARIS model of type **Table**.

In ARIS Architect, you can export the content of a model of the **Table** type and generate a link to the export file in the form of a URL. For more information about on how to create an ARIS table see the ARIS Architect online help. See also **Generate dashboard link** in the ARIS Architect online help.

The following parameters are available.

Parameters	Description
Source	<p>ARIS export file in XML format.</p> <p>Size limit: Unlimited.</p> <ul style="list-style-type: none"> ▪ URL: HTTP address of the source file If another operator supplies the URL dynamically, the URL cannot be edited here. Example <code>http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</code> ▪ URL alias: Alias of a URL configuration. Only URL aliases for that you have the Usage privilege are available. Select a URL alias. For details, see Manage URL aliases (page 198).
Insert parameter ()	Inserts user defined input parameters at the cursor position. The button is clickable only if at least one user input parameter, for example, Text user input , has been inserted in the feed definition.
Parameter options ()	Enables you to set input parameter options.
Refresh rate ()	Specifies the time until the data source is read in again. Default value is 12 h.
Authentication	Specifies an HTTP basic access authentication. User name and password are required for accessing the source file.
HTTP headers	Adds HTTP headers to the URL.
Parameters: Detect	Automatically reads out potential parameters of the data source. A requested parameters list is displayed that is based on the specified ARIS table source. You can enter the parameter values required to read in the data source.
Configure columns	Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset columns , the column list is reloaded from the data source and all changes in the list are undone.

9.2.1.2 BigMemory

Extracts data from a BigMemory cache.

To use the BigMemory source operator, you have to configure a BigMemory connection in MashZone NextGen in advance, see BigMemory for Caching, Connections and In-Memory Stores (page 167).

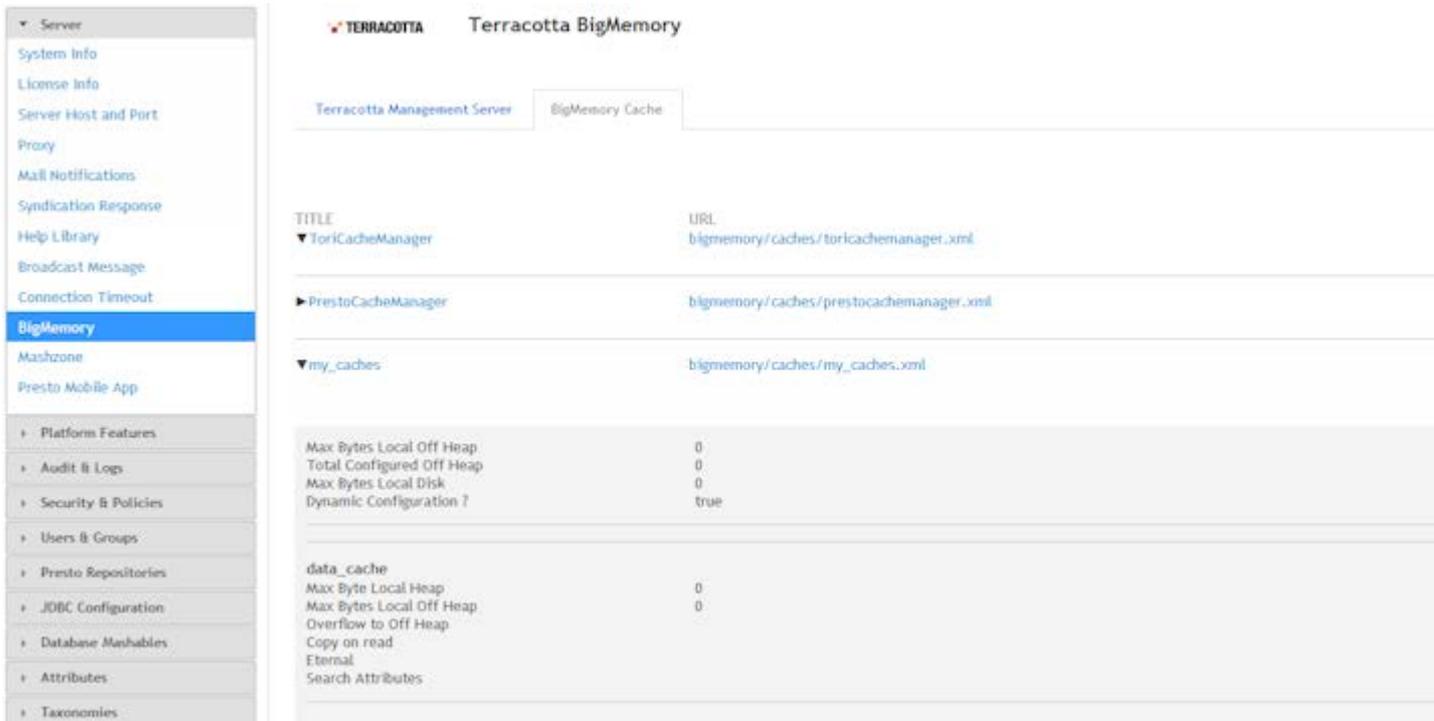
The following parameters are available.

Parameters	Description
Cache alias	List of available cache managers. Only Cache aliases for that you have the Usage privilege are available. Select an Cache alias.
Refresh rate ()	Specifies the time until the data source is read in again. Default value is 12 h.
RAQL query	Input field to enter any RAQL statement. Preset query: SELECT * FROM. See example BigMemory data source below.
Insert parameter ()	Inserts user defined input parameters at cursor position. The button is only clickable if at least one user input parameter, for example, Text user input , has already been inserted in the feed definition. See RAQL and SQL statement parameters (Seite 316) for details.
Expand ()/ Collapse ()	Expands or collapses the query input box.
Configure columns	Configures the columns list. Unselect a column for excluding it from the result data. Entering a New name for a column will cause that to be used instead of the original column name in the result. Clicking Reset columns will reload the column list from the data source and undo all changes in the list.

Example: BigMemory data source

To use the RAQL source operator, you have to configure a BigMemory connection in MashZone NextGen in advance, see BigMemory for Caching, Connections and MashZone NextGenMashZone NextGen (page 167).

To use the RAQL BigMemory operator you have to follow these conventions. The **Cache alias** pattern is <configuration name>.<cache name>. <cache name>. <configuration name> is the **Title** you specified as name for the configuration, <cache name> is one of the caches available in the configuration. In this example, the **Cache alias** is "my_caches.data_cache".



In the statement field, you can enter the RAQL statement. If you refer to a cache within the configuration, you just specify the cache name (i. e. without the configuration name). In the above example, if you wanted to address the "data_cache" cache, you would formulate your RAQL statement similar to this:

SELECT description FROM data_cache

In the case of a cache name that is incompatible with MashZone NextGen cache naming conventions, an alternative cache name is assigned in the list of available caches (**Cache**). To access such a cache, the alternative name (**Cache**) must be used in RAQL queries.

9.2.1.3 CSV

Reads the CSV file and writes the individual values (character strings) to the table columns in the data feed based on the specified parameters. A change of column is identified by the specified separator between the individual values.

The following parameters are available.

Parameters	Description
Source	Text file, with values that are separated by the same separator. <ul style="list-style-type: none"> ▪ URL: HTTP address of the source file If another operator supplies the URL dynamically, the URL

Parameters	Description
	<p>cannot be edited here.</p> <p>Example</p> <p><code>http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</code></p> <ul style="list-style-type: none"> ▪ URL alias: Alias of a URL configuration. Only URL aliases for that you have the Usage privilege are available. Select a URL alias. For details, see Manage URL aliases (page 198). ▪ Local file: Loads file from a resource directory. Files must be located in a defined resource directory on the MashZone NextGen server. <p>Path prefix (alias): Alias of the resource directory with the path to a directory on the server. Select an alias of the local resource directory.</p> <p>In the input box, enter a path to the relevant data source file for the Path prefix (alias) selected.</p> <p>Example</p> <p><code>/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</code></p>
Insert parameter 	<p>Inserts user defined input parameters at the cursor position. The button is clickable only if at least one user input parameter, for example, Text user input, has been inserted in the feed definition. For details, see URL parameter syntax (page 315).</p>
Parameter options 	<p>Enables you to set input parameter options.</p>
Refresh rate ()	<p>Specifies the time until the data source is read in again. Default value is 12 h.</p>
Authentication	<p>Specifies a HTTP basic access authentication or an existing authentication defined in the MashZone NextGen administration. User name and password are required for accessing the source file.</p>
HTTP headers	<p>Adds HTTP headers to the URL.</p>
Separator	<p>Separates the column values in the CSV file. Comma (,), semicolon (;), space, tab and pipe () are available. Default is comma.</p>
Get column names from row	<p>Specifies a specific row that contains the column names.</p> <p>Activate the Get column names from row option and enter the number of the relevant row.</p>
Data from row	<p>Specifies a specific row from which the data source values extraction starts.</p>

Parameters	Description
Advanced parsing options	<p>Charset: Character set in which the source file is coded. The default value can be changed manually if the extracted data refers to a different coding type. Default: windows-1252.</p> <p>Masking: Protects the enclosed characters against being split at the separator. If column values contain the specified separator, they can be enclosed in a pair of masking characters, for example, "1,23". Masking characters can be set as required (available masking characters: single-quote ('), double-quote ("), none).</p> <p>Remove quote characters in column value: Removes the characters used for masking from the result data. If this option is deactivated the masking characters will remain as part of the result data.</p> <p>Sanitize names: Transforms column names in such a way that they can be used as XML names. This affects the names containing blank spaces or other special characters. If this option is deactivated, the names are left unchanged.</p> <p>Trim whitespace: Removes all leading and trailing whitespace from column values. If this option is deactivated, whitespace can be part of the result data.</p>
Configure columns	<p>Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset columns, the column list is reloaded from the data source and all changes in the list are undone.</p>

Upload CSV files to the MashZone NextGen Repository. If required, you can upload CSV files to the MashZone NextGen Repository. For details, see Add External Resources as MashZone NextGen Files (page 195). To upload a file to the MashZone NextGen Repository administration privileges are required.

If you use an absolute URL, for example,

http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication must be set. All requests to the uploaded file are made with the specified user/password combination.

If you use a relative URL, for example,

/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication is not required. All requests to the uploaded file are made with the current logged in user.

In some cases, it might be necessary to grant view permission for a user. This can be done with the API console, e.g.:

```
{
  "version": "1.1",
  "sid": "PolicyService",
```

```
"svcVersion": "0.1",
"oid": "addPermissions",
"params": [
"fileName", "type.entity.file", "VIEW",
[{"principalId": "userID", "principalTypeId": "User"}]]
}
```

- **fileName** is the name that was specified while uploading the file.
- **userID** is the id of the user, who should get the permission.

9.2.1.4 Data feed

Extracts data from an existing data feed.

To include a data feed as data source in a dashboard or in other data feeds, the **View** permission for the data feed selected is required. See Manage data feed permissions (page 84) for detail.

The following parameters are available.

Parameters	Description
Data feed	Data feed selected. Select data feed displays a list of available data feeds.
Edit data feed	Opens the selected data feed in a new tab.
User inputs	List of user inputs used in the selected data feed. You can enter a value in the relevant input box.
Configure columns	Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset columns , the column list is reloaded from the data source and all changes in the list are undone.

9.2.1.5 Events

Quickly determines event query data from webMethods Events. Data is retrieved via a predefined Event Bus using a real-time buffer server.

The following parameters are available.

Parameters	Description
Event source	Reads data from a real-time buffer that can be selected from a list of configured real-time buffer instances. For details, see Event Service Configuration and Administration (Seite 221).
Configure columns	Configures the columns list. Unselect a column for excluding it from the result data. Entering a New name for a column will cause that to be used instead of

Parameters	Description
	the original column name in the result. Clicking Reset columns will reload the column list from the data source and undo all changes in the list.

9.2.1.6 JDBC

Extracts data from a preconfigured JDBC database.

You can configure JDBC data sources in the MashZone NextGen Administration. For details, see [Manage data sources and drivers \(page 180\)](#).

The following parameters are available.

Parameters	Description
Data source	List of available JDBC data sources. Only JDBC data sources for which you have the required privilege are available. Select a JDBC data source.
Refresh rate ()	Specifies the time until the data source is read in again. Default value is 12 h.
SQL query	Input field to enter any SQL query command, for example, <code>SELECT * FROM <table name></code> . Technically, execution of DDL (data definition language) and DML (data manipulation language) commands, such as create or insert is not prohibited, however, we recommend that you do not run this type of command with MashZone NextGen. To minimize the risk of accidental changes to the database schema used, we recommend accessing the database schema via a role or user with only a read-only privilege.
Insert parameter ()	Inserts user defined input parameters at the cursor position. The button is clickable only if at least one user input parameter, for example, Text user input , has been inserted in the feed definition. For details, see RAQL and SQL statement parameters (Seite 316).
Expand ()/ Collapse ()	Expands or collapses the query input box.
Authentication	Specifies an HTTP basic access authentication. User name and password are required for accessing the source file.
Configure columns	Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset

Parameters	Description
	columns , the column list is reloaded from the data source and all changes in the list are undone.

9.2.1.7 JSON

Extracts data from a JSON file.

The following parameters are available.

Parameters	Description
Source	<p>JSON file</p> <ul style="list-style-type: none"> URL: HTTP address of the source file If another operator supplies the URL dynamically, the URL cannot be edited here. Example http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1 URL alias: Alias of a URL configuration. Only URL aliases for that you have the Usage privilege are available. Select a URL alias. For details, see Manage URL aliases (page 198). Local file: Loads file from a resource directory. Files must be located in a defined resource directory on the MashZone NextGen server. Path prefix (alias): Alias of the resource directory with the path to a directory on the server. Select an alias of the local resource directory. In the input box, enter a path to the relevant data source file for the Path prefix (alias) selected. Example /documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1
Browse file alias (...)	<p>Enables you to browse the resource directories with the alias defined. Click the Browse file alias (...) button and select the required source file. At least one resource directory must exist. For details, see Manage resource directories (page 196).</p>

Parameters	Description
Insert parameter 	Inserts user defined input parameters at the cursor position. The button is only clickable if at least one user input parameter, for example, Text user input , has been inserted in the feed definition. For details, see URL parameter syntax (page 315).
Parameter options 	Enables you to set input parameter options
Refresh rate ()	Specifies the time span before the source file is read in again. Default valueThe default value is 30 sec.
Authentication	Specifies a HTTP basic access authentication or an existing authentication defined in the administration. User name and password are required for accessing the source file.
HTTP headers	Adds HTTP headers to the URL
Repeating object: Detect	<p>Reads out the repeating object automatically. A repeating object already set by a user will not be considered. The detected columns are displayed in the Columns box depending on the repeating object.</p> <p>JSON object that is repeated for each row; You can edit the repeating object manually.</p> <p>Specify a valid XQuery 3.1 lookup expression, for example,</p> <pre>?catalog?journal?articles?* ?data?rows?* ?catalog?books?*</pre> <p>A JSON document is loaded as a nested structure of maps and arrays. The above expressions use the (terse form of the) map/array lookup syntax as follows:</p> <ul style="list-style-type: none"> - The initial '?' is a unary lookup operator that selects a named member of the root map - Further '?' are postfix lookup operators - '?*' selects all members of an array, for example, '?5' selects the fifth element of an array. <p>For more information about the lookup syntax, see https://www.w3.org/TR/xquery-31/#id-lookup.</p>
Advanced parsing options	<p>Charset: Character set in which the source file is coded. This can be set manually if the extracted data refers to a different coding type.</p> <p>By default, the HTTP response encoding is used if available, otherwise UTF-8 is used.</p>
Columns: Detect	Reads out the columns of the data source automatically. The requested columns lists are displayed, based on the specified

Parameters	Description
	repeating object.
Configure columns	Configures the columns list. Deselect a column to exclude it from the result data. If you enter a New name for a column, it is used instead of the original column name in the result. If you click Reset columns , the column list is reloaded from the data source and all changes in the list are undone.

Upload JSON files to the MashZone NextGen Repository. If required, you can upload CSV files to the MashZone NextGen Repository. For details, see Add External Resources as MashZone NextGen Files (page 195).

To upload a file to the MashZone NextGen Repository administration privileges are required.

If you use an absolute URL, for example,

http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication has to be set. All requests to the uploaded file are made with the specified user/password combination.

If you use a relative URL, for example,

/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication is not required. All requests to the uploaded file are made with the current logged in user.

In some cases, it might be necessary to grant view permission for a user. This can be done with the API console, e.g.:

```
{
  "version": "1.1",
  "sid": "PolicyService",
  "svcVersion": "0.1",
  "oid": "addPermissions",
  "params": [
    "fileName", "type.entity.file", "VIEW",
    [{"principalId": "userID", "principalTypeId": "User"}]]
}
```

- **fileName** is the name that was specified while uploading the file.
- **userID** is the id of the user, who should get the permission.

9.2.1.8 Excel

Reads a worksheet of an MS Excel file and writes the individual values to table columns in the data feed based on the specified parameters. The source table can be imported as a list or cross table.

In list tables, a corresponding column is created in the data feed for every non-empty column in the source table.

In cross tables, three columns are created in the data feed: A vertical iteration column corresponding to the first source column with the header, a horizontal iteration column defined in the operator and a value column.

MS Excel cells of data type number are extracted accurately, regardless of their formatting.

Therefore, the values can be more accurate than displayed in MS Excel. MS Excel cells of data

type date are extracted according to the formatting information to maintain the accuracy of the time stamp.

A cell can have a maximum of 2,000 characters.

The following parameters are available.

Parameters	Description
Source	<p>MS Excel file (xls, xlsx)</p> <ul style="list-style-type: none"> ▪ URL: HTTP address of the source file If another operator supplies the URL dynamically, the URL cannot be edited here. Example <code>http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</code> ▪ URL alias: Alias of a URL configuration. Only URL aliases for that you have the Usage privilege are available. Select a URL alias. For details, see Manage URL aliases (page 198). ▪ Local file: Loads file from a resource directory. Files must be located in a defined resource directory on the MashZone NextGen server. Path prefix (alias): Alias of the resource directory with the path to a directory on the server. Select an alias of the local resource directory. In the input box, enter a path to the relevant data source file for the Path prefix (alias) selected. Example <code>/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</code>
Browse file (alias)	<p>Enables you to browse the resource directories with the alias defined. Click the Browse file alias (...) button and select the required source file. At least one resource directory must exist. For details, see Manage resource directories (page 196).</p>
Insert parameter ()	<p>Inserts user defined input parameters at the cursor position. The button is clickable only if at least one user input parameter, for example, Text user input, has been inserted in the feed definition. For details, see URL parameter syntax (page 315).</p>
Parameter options ()	<p>Enables you to set input parameter options.</p>
Refresh rate ()	<p>Specifies the time until the data source is read in again. Default value is 12 h.</p>

Parameters	Description
Sheet	<p>Sheet in the source table to be extracted.</p> <p>The default value: First sheet</p> <p>Specification: Mandatory</p>
List table / Cross table	<p>Specifies the table type.</p> <p>The default value: List table</p> <p>Specification: Mandatory</p> <p>For cross tables, only a single vertical iteration on the left side of the table is currently supported.</p>
Separator	<p>Separates the column values in the CSV file. Comma (,), semicolon (;), space, tab and pipe () are available. Default is comma.</p>
Column name from row	<p>Determines the names of the individual columns from a specific row, the row number of which must be specified.</p> <p>This option is not available for cross tables.</p>
Horizontal iteration from row	<p>Determines the column names of the individual iteration steps from a specific row. The column name of the vertical iteration is also determined from this row.</p>
Import values from row	<p>Extracts all values from the source file starting with a specific row.</p> <p>The default value: 2</p> <p>Specification: Mandatory</p>
Import data range from/to	<p>Area of the table from which data is to be extracted, specified using column and row coordinates, for example, A3 to H128</p> <p>Specification: Optional</p> <p>You can only specify a single continuous data range that can contain empty rows or columns.</p> <p>If no upper limit (to) is specified for the data range, all cells above the lower limit (from) are extracted.</p>
Configure columns	<p>Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset columns, the column list is reloaded from the data source and all changes in the list are undone.</p>

Upload XLS files to the MashZone NextGen Repository If required, you can upload XLS files to the MashZone NextGen Repository.

For details, see Add External Resources as MashZone NextGen Files (page 195).

To upload a file to the MashZone NextGen Repository administration privileges are required.

If you use an absolute URL, for example,

http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37

bc9fa44580, an authentication must be set. All requests to the uploaded file are made with the specified user/password combination.

If you use a relative URL, for example,

/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication is not required. All requests to the uploaded file are made with the current logged in user.

In some cases, it might be necessary to grant view permission for a user. This can be done with the API console, e.g.:

```
{
  "version": "1.1",
  "sid": "PolicyService",
  "svcVersion": "0.1",
  "oid": "addPermissions",
  "params": [
    "fileName", "type.entity.file", "VIEW",
    [{"principalId": "userID", "principalTypeId": "User"}]]
}
```

- **fileName** is the name that was specified while uploading the file.
- **userID** is the id of the user, who should get the permission.

9.2.1.9 PPM

Uses the Process Performance Manager (PPM) query interface to retrieve data from favorites defined in PPM.

For every PPM system that can be used in ARIS Connect, a PPM connection must be created.

The relevant PPM client server must be running. See the PPM documentation **PPM Installation** for details.

In PPM, you can use the pop-up menu of a favorite to copy the corresponding favorite URL. Click **Use in dashboard** in the pop-up menu of a favorite. See the PPM online documentation for details.

If you have updated ARIS Connect 10.0 Service Release 1 to version 10.0 Service Release 2, you must recreate the PPM connections that were defined for ARIS Connect 10.0 Service Release 1. First delete the PPM connection and add a new PPM connection with the same data.

The following parameters are available.

Parameters	Description
PPM Connection	Alias of the PPM Connection, that contains the PPM client connection data defined in ARIS Connect. See Configure PPM server for details.
Refresh rate ()	Specifies the time until the data source is read in again. Default value is 12 h.
Favorite	Path of a PPM favorite. The favorite path represents the favorites tree including favorites folder and name, for example, \Favorites\Process cycle time
Extract from URL	Determines the connection data of the PPM data source automatically, for example, alias, favorite path, language and favorite type. Click Extract from

Parameters	Description
	URL and insert the favorite URL created in PPM.
Authentication	<p>Specifies the credentials for authenticating the query against the PPM client server.</p> <p>Single Sign-On: Enables you to log in to PPM client server using single sign-on (SSO), and your current credentials. For details on how to configure SSO, see Configure single sign-on.</p> <p>HTTP basic auth: Requires the user name and the password of a PPM user. The returned data is filtered based on the PPM user access rights. The PPM user must have access rights for the selected favorite.</p>
Advanced options	<p>Specifies further connection parameters.</p> <p>Favorite type specifies the favorite as Private or Shared.</p> <p>Language of the favorite, for example, en for English</p> <p>Request key columns separately: Extracts all key values from the list table and writes them to separate columns of the data feed.</p>
Configure columns	Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset columns , the column list is reloaded from the data source and all changes in the list are undone.
Filters: Detect	Reloads the dimensions and measures information from PPM
Add filter	Adds one or more filter criteria to filter the values of the PPM query. See Add filter (page 310) for details.

9.2.1.9.1 Add filter

You can add one or multiple filter criteria to filter the values of the PPM query.

Procedure

1. Click **Add filter**.
2. Select a **Dimension** or **Measure** provided by PPM, for example., date or process throughput time, and click **Add**. The selected criterion is added to the PPM operator.
3. Select an operator for the condition of the filter criterion in the drop-down menu, for example, is equal to.
4. Enter a constant filter value in the input field or insert an input parameter by clicking  **Insert input parameter**. The button is only clickable if there is at least one input parameter with the same data type.
5. Click the **Set filter properties icon** and specify the criterion parameters.

When specifying parameters, such as filter criterion and filter values, follow the predefined syntax of the PPM query interface (query API). See example below. For detailed information on the PPM query interface, refer to the PPM documentation **PPM Query interface**.

Filter using an expression uses criterion expressions for filtering. This option is available for criterions of data type **text**.

Key uses the criterion key as a filter value.

Description uses the criterion description as a filter value.

Set filter level allows the selection of the level you want to filter. This option is available for two level and multi level dimensions.

Example: The two level dimension **Dealer** with the levels **Region** and **City** allows you to filter the data according to Region or City.

Scaling. This option is available for criterions of data type **date** and for some numeric data types (for example, a numeric or percentage KPI value).

6. Click  **Add condition** to add further filter conditions.

7. Click **Next** to close the dialog.

Your settings are applied.

Example

The result of the PPM query depends on the syntax used in the PPM query interface. The character combination of space + (in the filter value (for example, ABC (123)) can lead to different results, depending on the filter criterion.

If you filter for "SOLDTO_NAME=ABC (123)", PPM searches for the customer with the name = "ABC" and the description = "123".

If you filter for "SOLDTO_NAME(VAL)=ABC (123)", PPM searches for the customer name "ABC (123)" and does not take the description into account.

If you filter for "SOLDTO_NAME(DESC)=ABC (123)", PPM searches for the customer whose description is "ABC (123)" and does not take the name into account.

9.2.1.10 Terracotta DB

Extracts data from datasets managed by a Terracotta DB server and runs RAQL queries over the data. The available datasets were either generated by an external application or by MashZone NextGen for materializing data feed results (page 86). You must select one of those two categories of datasets. Depending on your selection, the source operator shows different parameters.

To use the Terracotta DB source operator, you must configure a Terracotta DB connection in MashZone NextGen in advance. Terracotta DB connections can be defined in the MashZone NextGen **Admin console**. See Manage Terracotta DB connections (Seite 177).

The following parameters are available for the **Datasets** and **Materialized Feeds** source types.

Parameters	Description
Source	<p>Specifies the data source type.</p> <ul style="list-style-type: none"> ▪ Datasets ▪ Materialized Feeds <p>Select a source type.</p>
Terracotta DB alias	<p>List of available Terracotta DB datasets for the configured connection aliases.</p> <p>Only Dataset aliases for that you have the Usage privilege are available. Select a Dataset alias. See Manage Terracotta DB connections (Seite 177) for details.</p> <p>The option is available if the Datasets source type is selected.</p>
Materialized Feed	<p>List of available Materialized Feed aliases.</p> <p>Only Materialized Feed aliases for that you have the Usage privilege are available. Select the alias which defines the dataset with the Materialized Feed results. The corresponding dataset is available in the Dataset name drop-down menu. See Manage Materialized Feeds (page 188) for details.</p> <p>The option is available if the Materialized Feeds source type is selected.</p>
Dataset name	<p>List of available Terracotta DB datasets for the configured connection alias. Select a dataset.</p> <p>The option is available if the Datasets source type is selected.</p>
Refresh rate ()	<p>Specifies the time until the data source is read in again. Default value is 12 h.</p>
RAQL query	<p>Input field to enter any RAQL statement.</p> <p>Preset query: <code>SELECT * FROM Dataset</code></p> <p>Dataset is the selected entry from the Dataset Alias drop-down box.</p>
Insert parameter ()	<p>Inserts user defined input parameters at the cursor position. The button is clickable only if at least one user input parameter, for example, Text user input, has been inserted in the feed definition. For details, see RAQL and SQL statement parameters (Seite 316).</p>
Expand ()/ Collapse ()	<p>Expands or collapses the query input box.</p>
Configure columns	<p>Configures the columns of the dataset.</p> <p>The option is available if the Datasets source type is selected.</p> <p>When selecting a dataset, schema guessing is used to derive its</p>

Parameters	Description
	<p>columns. Using 1000 records, the guessing mechanism derives the set of all unique cell definitions of those records and defines corresponding columns. Use Configure columns to display or modify that schema.</p> <p>Unselect a column to exclude it from processing. Select or enter a (positive) Sample size and click Derive columns to re-run the schema guessing with a larger sample size, thereby omitting the current configuration. Note that the larger the sample size is, the longer the schema guessing takes.</p> <p>Click Add column to manually add a new column. Enter a unique column name and select a data type for the new column. Note that the combination of name and type has to fit to the cell definition in the store. If not, the corresponding values will always be null. For example, if the cell definition is age with type INT and you specify age with type DOUBLE, the column definition does not fit to the cell definition and the values are null.</p>

9.2.1.11 XML

Extracts data from an XML file. The data records are identified using a recurring element. The individual values are written to the table columns in the data feed based on the specified parameters.

The following parameters are available.

Parameters	Description
Source	<p>XML file</p> <ul style="list-style-type: none"> ▪ URL: HTTP address of the source file If another operator supplies the URL dynamically, the URL cannot be edited here. Example <code>http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</code> ▪ URL alias: Loads the file from a resource directory. Files must be located in a defined resource directory on the server. The data source files are located in a folder or a sub-folder of the ARIS document storage accessible in the ARIS Connect Repository. If you have updated ARIS Connect 10.0 Service Release 1 to version 10.0 Service Release 2, you need to recreate the URL alias that have been defined for ARIS Connect 10.0 Service Release 1.

Parameters	Description
	<p>Delete the URL alias at first and add a new URL alias with the same data.</p> <p>Path prefix (alias): Alias of the resource directory with the path to a directory on the server. Select an alias of the local resource directory, for example, Feed URL.</p> <p>Feed URL is the alias of the resource directory of the default data feeds provided in ARIS Connect.</p> <p>In the input box, enter a path to the relevant data source file for the Path prefix (alias) selected.</p> <p>Example</p> <p>/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580?tenantid=default&revision=1</p> <p>To configure a Path prefix (alias), see Configure Dashboard server for details.</p> <p>To store data source files in a resource directory, see Upload file based data sources for details.</p> <p>To copy the source file URL, see Upload file based data sources for details.</p>
Insert parameter ()	<p>Inserts user defined input parameters at the cursor position. The button is clickable only if at least one user input parameter, for example, Text user input, has been inserted in the feed definition.</p>
Parameter options ()	<p>Enables you to set input parameter options.</p>
Refresh rate ()	<p>Specifies the time until the data source is read in again. Default value is 12 h.</p>
Authentication	<p>Specifies an HTTP basic access authentication.</p> <p>User name and password are required for accessing the source file.</p>
HTTP headers	<p>Adds HTTP headers to the URL.</p>
Repeating element: Detect	<p>Reads out the repeating element automatically. A repeating element already set by a user is not considered. The repeating element is displayed corresponding to the hierarchy of the XML elements, for example, <element>/<repeating element>. The requested columns are displayed in the Columns box depending on the repeating element.</p> <p>XML element that is repeated for each row (XPath to repeat element);</p> <p>You can edit the repeating element manually.</p> <p>Specify a valid XPath expression, for example,</p> <p>/catalog/journal/article</p>

Parameters	Description
	/data/row /catalog/book
Columns: Detect	Automatically reads out the columns of the data source. The requested columns list is displayed, based on the specified repeating element.
Configure columns	Configures the columns list. You can unselect a column to exclude it from the result data. If you enter a new name for a column, this name is used in the result instead of the original column name. If you click Reset columns , the column list is reloaded from the data source and all changes in the list are undone.

If you use an absolute URL, for example,

http://myhost.company:1080/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication has to be set. All requests to the uploaded file are made with the specified user/password combination.

If you use a relative URL, for example,

/documents/rest/links/03388871-367d-4abe-a2c9-37bc9fa44580, an authentication is not required. All requests to the uploaded file are made with the current logged in user.

9.2.1.12 URL parameter syntax

You can insert parameters in the input URL of the **ARIS Table**, **CSV**, **Excel**, **JSON**, and **XML** operators, as well as in the configuration of a **Call URL** action.

A parameter reference is represented by the parameter name in square brackets. If the name contains special characters, it must be enclosed in quotes.

Examples

http://[domain]

["web address"]

To represent a literal square bracket, you double the square bracket. In this way, the square bracket does not introduce a parameter reference. For example, **http://[::1]:8080** does not refer to a parameter, but represents the IPv6 loopback address literal **[::1]**.

At runtime, parameter references will be replaced by the value of their parameter.

A parameter reference can have options appended to it. Options appear inside of the square brackets and are separated by commas. The following options are supported:

- **encode-value="true"** causes URL encoding of the resulting parameter string
- **format="format string"** causes a numeric or date value to be converted to text with the given format

Examples

http://localhost:8080/[path,encode-value="true"]

http://localhost:[port,format="1234"]

9.2.1.13 RAQL and SQL statement parameters

You can insert parameter references in the RAQL and SQL statements of the **BigMemory**, **JDBC**, **Terracotta DB**, and **RAQL Inline** operators. They are represented by a colon, followed by the parameter name. If the name contains special characters, it must be enclosed in double quotes.

9.2.1.13.1 JDBC parameters

Parameters are handed over to the database server as typed values along with the SQL statement at execution time. Before a statement is handed over to the database server, it is rewritten to use JDBC parameter markers ('?') as follows:

- A single-value parameter is replaced by a single JDBC parameter marker.
- A reference of a non-empty list is replaced by a comma-separated list of JDBC parameter markers, one for each list element. This is useful for populating the values of an **IN** predicate by a list value.

Example

```
SELECT * FROM table WHERE column IN (:list)
```

with 3 elements in the list is rewritten to

```
SELECT * FROM table WHERE column IN (?,?,?)
```

- A reference of an empty list is replaced by a single parameter marker that refers to a **NULL** value.

The last rule provides a valid SQL statement for an empty list. Note that if the list of values is empty, both **IN** and **NOT-IN** predicates return the same result: UNKNOWN and not FALSE.

9.2.1.13.2 RAQL parameters

Parameters are handed over to the RAQL engine as typed values along with the RAQL statement at execution time. RAQL statements are usually passed unchanged to the RAQL engine. If they contain list parameters, they are rewritten to use single value parameters as follows:

- A reference of a non-empty list is replaced by a comma-separated list of parameter references, one for each list element. This is useful for populating the values of an **IN** predicate by a list value.

Example

```
SELECT * FROM feed1 WHERE column IN (:list)
```

with 3 elements in the list is rewritten to

```
SELECT * FROM feed1 WHERE column IN (:"list[1]","list[2]","list[3]")
```

In this case, "list[1]" etc. are generated names for the individual list elements.

- A reference of an empty list is not replaced, but a single **NULL** value is passed to the RAQL engine.

The last rule provides a valid RAQL statement for an empty list. Note that if the list of values is empty, both predicates **IN** and **NOT-IN** return the same result UNKNOWN and not FALSE.

9.2.2 Data transformation operators

In addition to the data source operators you can add further operators to transform the source data. The following data transformation operators are available in the data feed editor.

9.2.2.1 Aggregate

Combines rows if identical values occur multiple times in specified dimension columns. The numerical values in the remaining columns are combined using Average, Sum, Minimum, Maximum, or Number.

One or more numerical columns in a table are aggregated using none, one, or several dimension columns. In all rows with identical values in all dimension columns, the values in the columns to be aggregated are combined into one row based on an aggregation rule. That is, the result contains one row for each combination of dimension columns. This also applies if no columns are specified for aggregation. If no dimension column is specified, only a single row is created and all values in the columns to be aggregated are combined into a single value for each column. No rows are created for combinations that do not occur in the original table.

9.2.2.1.1 Parameters

The following parameters are available.

Action	Result
Dimension column	<p>Name of dimension column.</p> <p>Source: Source table</p> <p>Data type: Date, Number, or Text</p> <p>The default value: {None}</p> <p>Specification: Optional</p>
Aggregation column	<p>Name of the column to be aggregated.</p> <p>Source: Aggregation column is transferred from the source table.</p> <p>Data type: Number</p> <p>The default value: {None}</p> <p>Specification: Optional</p>
Aggregation type	<p>Aggregation type for the column to be aggregated: Average value, Sum, Minimum, Maximum, Count, First row and Last row.</p>

Action	Result
	The default value: Average value, if aggregation column selected. Specification: Mandatory, if aggregation column selected.
Weighting	If the aggregation type is Average value or Sum, a numerical column for weighting of the rows can be specified for each column to be aggregated. Specification: Optional

9.2.2.1.2 Aggregation type

The following aggregation operations can be applied to the columns to be aggregated.

Data type	Comparison operators
Minimum	Finds all rows that have a specific combination of values in the dimension columns and returns the lowest value that occurs in these rows in the column to be aggregated.
Maximum	Finds all rows that have a specific combination of values in the dimension columns and returns the highest value that occurs in these rows in the column to be aggregated.
Average value	Finds all rows that have a specific combination of values in the dimension columns and returns the average of the values in the column to be aggregated. For weighting purposes, an additional column can be specified for each source column, containing a weighting factor for each row. The weighting information is combined as a pair with the source column.
Sum	Finds all rows that have a specific combination of values in the dimension columns and returns the sum of the values in the column to be aggregated. For weighting, an additional column can be specified for each source column, containing a weighting factor for each row. The weighting information is combined as a pair with the source column.
Count	Finds all rows that have a specific combination of values in the dimension columns and returns the count of values in the column to be aggregated.
First row	Finds all rows that have a specific combination of values in the dimension columns and returns the value of the row with the lowest row index (according to the index column).
Last row	Finds all rows that have a specific combination of values in the dimension columns and returns the value of the row with the highest row index (according to the index column).

At least one dimension or aggregation column, or both, must be set.

If no aggregation columns or dimension columns are specified, the incoming table remains unchanged.

9.2.2.1.3 Examples

The following table is to be aggregated based on the **Dim 1** and **Dim 2** columns. The sum is to be calculated for the Values 1 column and the average for the Values 2 column. The Weight (values 2) column is used for weighting the Values 2 column one row at a time.

Dim 1	Dim 2	Values 1	Values 2	Weight (values 2)
A	X	1	2	3
B	Y	2	4	4
C	Z	5	6	3
A	X	7	8	4
B	Y	9	10	3
C	Z	11	12	4

Result

Dim 1	Dim 2	Sum (values 1)	Average (values 2)
A	X	8 (1+7)	5,43 (2*3 + 8*4)/(3+4)
B	Y	12 (3+9)	6,57 (4*4 + 10*3)/(4+3)
C	Z	16 (5+11)	9,43 (6*3 + 12*4)/(3+4)

9.2.2.2 Arithmetic

Executes various arithmetical calculations. The operator sets any number of numerical operands against each other. The values are set against each other one row at a time according to the specified calculation type. The calculation always runs from top to bottom. In other words, two operands are always set against each other one row at a time and the result from the first two operands is then set against the third operand.

Compounding can be mapped by using a separate operator for each expression in brackets.

9.2.2.2.1 Parameters

The following parameters are available.

Action	Result
Operands	<p>One numerical operand per operation for the "Square" and "Square root" calculation types, otherwise two numerical operands.</p> <p>Source: Source table, constants, user input or incoming values from other operators.</p> <p>Data type: Number</p> <p>Specification: Mandatory</p>
Calculation type	
Addition (+)	Adds two columns row by row
Subtraction (-)	Subtracts 2 columns row by row
Division (/)	Divides the first column by the second column
Multiplication (*)	Multiplies two columns row by row
Percent (%)	Row by row, multiplies the second column with the percent value of the first column
Square (x ²)	Calculates the square of a column
Root	Calculates the square root of a column
Sine (sin)	Calculates for a column the sine value of an angle in degrees
Cosine (cos)	Calculates for a column the cosine value of an angle in degrees
Tangent (tan)	Calculates for a column the tangent value of an angle in degrees
Arcsine (asin)	Calculates for a column the arcsine of an angle in degrees
Arccosine (acos)	Calculates for a column the arccosine of an angle in degrees
Arctangent (atan)	Calculates for a column the arctangent of an angle in degrees
Logarithm (lg)	Calculates for a column the common logarithm
Logarithm (ln)	Calculates for a column the natural logarithm
Power (exp)	Calculates for a column S1 the S2nd power of S1 (S1 to the power of S2)
Minimum (min)	Calculates the minimum of column 1 and column 2
Maximum (max)	Calculates the maximum of column 1 and column 2
Absolute value (abs)	<p>Calculates for a column the absolute value</p> <p>The default value: Addition (+)</p> <p>Specification: Mandatory</p> <p>You need to specify the source values for the trigonometric functions sin,</p>

Action	Result
	cos, tan, asin, acos, and atan in degrees.

9.2.2.2.2 Example

Example: Result = Column 1 + Column 2 - Column 3

Column 1	Column 2	Column 3	Result
1000	2000	50	2950
2000	3000	1000	4000
3000	4000	1200	5800

9.2.2.3 Average

Calculates the average of the values from several numerical source columns one row at a time, writes the result to a target column, and overwrites any existing values there. If the target column does not exist, it is created.

9.2.2.3.1 Parameters

The following parameters are available.

Parameter	Description
Column	Name of the column for which the average value is calculated. Column is transferred from the source table. Data type: Number Specification: Mandatory
Weight	Weight factor, which can be specified for each column to be aggregated: a column with values, a single value from a feed (single-value operator), an input value or a constant. Data type: Number Specification: Optional
Target column	Name of the column to which the result is written. The column name can be transferred from the source table or freely entered. Data type: Number Default value: Result_1

Parameter	Description
	Specification: Mandatory

9.2.2.4 Change data type

Changes the data types of the specified columns to the **Number**, **Text**, or **Date** data types

Action	Result
Conversion of Text to Number	<p>Numerical value of the text taking into account the decimal separator.</p> <p>If the decimal separator is set correctly, any thousands separator is detected automatically.</p>
Conversion of Number to Text	<p>Text representation of the number in the internal format, or based on the language and the specified format. You can also specify a valid number of leading zeros.</p> <p>If nothing is specified here, the results are formatted in the numerical format.</p>
Conversion of Text to Date	<p>Date value of the text in the internal format, based on the specified format and, where applicable, the language.</p> <p>The date must be in the AD era. Date values before the common era are not supported. The time format must be specified. The time format is made up of sequences of characters, which stand for date fields, for example, year, month, day of the week, or minute, in the relevant language; separated by separators. In addition, the corresponding language must be specified. Non-editable text must be enclosed in quotation marks.</p> <p>When using the Q or q symbol for quarters, all other symbols except Y and y for years are ignored. Only the order of Q/q and Y/y is relevant.</p>
Conversion of Date to Text	<p>Text representation of the date in the internal format, or based on the language and the specified format</p> <p>Non-editable text must be enclosed in quotation marks.</p> <p>The format and language specifications are optional. If no format is specified, the data is output in the internal date format. If no language is specified, English (EN) is applied as the default language.</p>
Conversion of Number to Date	<p>Date value corresponding to the value of the number as milliseconds since 01/01/1970</p>

Action	Result
Conversion of Date to Number	Number of milliseconds since 01/01/1970

Internal number format

If a user is logged in in English, the number format is Anglo-Saxon, which uses a period as the decimal separator, but no grouping characters.

Internal date format

Use yyyy-Q for specifying to the nearest quarter, otherwise use yyyy-MM-ddThh:mm:ss. The number of digits corresponds to the accuracy of the date, and the remaining digits are omitted. This is the transfer format.

Quarterly specifications

Quarterly specifications are indicated by a **Q** within the section of the format that is not in single quotation marks.

Prerequisites for conversion of text into quarterly date values:

1. It is expected that a source value containing a quarterly date consists of just two sequences of figures indicating the year and the quarter. Any non-numerical characters can occur before, after and between them, for example, Quarter 04/2009.
2. The pattern uses **Y** or **y** as the symbol for the year and **Q** or **q** for the quarter, for example, quarter Q/y or Q Y.

Procedure:

1. The (first) two sequences of digits are determined from the source value.
2. The section of the format that is not enclosed in between single quotation marks is used to determine whether **q / Q** or **y / Y** appears first.
3. If q or Q appears first, the first sequence of digits is interpreted as the quarter and the second as the year, and vice-versa.

9.2.2.4.1 Parameters

The following parameters are available.

Parameter	Description
Column	Name of the column to be changed; Source: Source table Data type: Date, Number, or Text Specification: Mandatory
New type	New column data type Default value: Text Specification: Mandatory

Parameter	Description
Format (Date type)	<p>Time format for conversion from Date type to Text type and vice versa.</p> <p>The following formatting symbols are available when converting date into text:</p> <p>Year: y or Y</p> <p>Quarter: Q</p> <p>Month: M</p> <p>Calendar week: w</p> <p>Day of the week: E or e</p> <p>Day of the month: d</p> <p>Day of the year: D</p> <p>Hour: H or h</p> <p>Minute: m</p> <p>Second: s</p> <p>AM/PM: a</p> <p>Time zone: z (for example, GMT)</p> <p>RFC time zone: Z (for example, -0900)</p> <p>Era: G (must always be AD)</p> <p>Default value: MM/dd/yyyy</p>
Format (Date type)	<p>The following formatting symbols are available when converting text to date:</p> <p>Year: y</p> <p>Quarter: Q</p> <p>Calendar week: w</p> <p>Week of the month: W</p> <p>Day of the week: E</p> <p>Day of the month: d</p> <p>Day of the year: D</p> <p>Hour (0-23): H</p> <p>Hour (1-24): k</p> <p>Hour AM/PM (0-11): K</p> <p>Hour AM/PM (1-12): h</p> <p>Minute: m</p> <p>Second: s</p> <p>AM/PM: a</p> <p>Time zone: z</p> <p>RFC time zone: Z</p>

Parameter	Description
	<p>Era: G</p> <p>Default value: yyyy-MM-dd'T'HH:mm:ss;</p> <p>Permitted separators in both cases:</p> <p>Dash/minus (-), underscore (_), slash (/), period (.), colon (:), comma (,), tab character, and space.</p> <p>Specification: Mandatory</p>
Language	<p>Language if the target format is of the Date type.</p> <p>Available languages: de and en.</p> <p>Specification: Mandatory when using names of months and names of days of the week</p>
Decimal separator	<p>Separator for the decimal places, if the target format is of the Number type.</p> <p>Default value: Comma (,)</p> <p>Specification: Mandatory</p>
Format (Number source format)	<p>Number format for the conversion of the Number type to the Text type. You can select predefined formats or set your own format manually.</p> <p>If the numbers before the decimal separator are entered manually, they must have four digits in ascending order followed by digits in descending order. This combination can be followed by text (such as the unit "hours" or km/h).</p> <p>Default value: 1234</p> <p>Permitted separators:</p> <p>Thousands separator in German: period (.)</p> <p>Thousands separator in English: comma (,)</p> <p>Decimal separator in German: comma (,)</p> <p>Decimal separator in English: period (.)</p>
Leading zeros	<p>Number of leading zeros. The maximum number of leading zeros is the number of digits before the decimal separator.</p> <p>Example:</p> <p>Format: 1,234.12 and leading zeros: 5</p> <p>Number -> Text</p> <p>10,245 -> 00010,25</p> <p>12000,4 -> 12000,4</p> <p>89,7 -> 00089,70</p>
Specify type	<p>Automatically specifies the data type of the source values.</p> <p>If the content of a column does not correspond to its assigned data type,</p>

Parameter	Description
	<p>a row is created in the operator which specifies the data type determined for this column. Vice-versa, settings (lines) are removed from the operator, which would reset the data type of a column already typified correctly.</p>
<p>Encoding (Text source format)</p>	<p>Specifies the encoding of special characters, for the conversion of the Text type to the Text type, for example, "/", "&", "?".</p> <p>Apply UTF-8 encoding: encodes the entire text, using UTF-8 codes</p> <p>Decode UTF-8: decodes the entire text, using UTF-8 codes</p> <p>Apply URL encoding: keeps the URL specific characters</p> <p>For UTF-8 encoding/decoding the special characters must be masked in a valid URL. Only select this option if you are sure that all sections that make up the URL are already masked.</p>

The characters used for the time format can be combined in any order and repeated any number of times.

Exceptions:

- For a month, the number of characters must be ≥ 3 (MMM or MMMM) if the month is specified in text format (JAN, FEB, etc.) and < 3 if it is specified as a figure. In this case, a language must also be specified so that the name of the month can be transformed correctly.
- For a year format, such as 2009, y can be specified any number of times, that is, yy and yyyy return 2009.
- For a year format, such as 09, however, yyyy returns the year 9 and yy the year 2009.
- When formatting date values as days of the week for a date-to-text conversion, an e/E number < 4 returns the day abbreviations (MON, TUE, etc.), while e/E = or > 4 returns the full name of the day.
- Only the month (M), minute (m), time zone (z), RFC time zone (Z) and calendar week (w) are case-sensitive.
- When converting text to date, if the Q or q symbol is used for quarters, all other symbols except Y and y are ignored. Only the order of Q/q and Y/y is key.
- For the reverse conversion from date to text, the Q/q can be combined with any other symbols, but may only occur once (not QQ/yy)
- Quarter entries are currently only possible in the form YYYY-Q. These strings can contain only the year, separator, and quarter.
- All other strings must be enclosed in single quotation marks ('). Spaces can be inside or outside, for example, 'On' dd.MM.yy 'at' hh:mm, or 'On 'dd.MM.yy' at 'hh:mm' '.
- The space pattern in the source and target format must match, for example, "2 .3 .09" -> "d .M .y" but not "2. 3. 09" -> "d .M .y".

9.2.2.4.2 Examples

Source format: "22.3.2009"

Time format: "d.M.y" or "DDDD.MM.YYYY",
but not "DD.MMM.YYYY "

Source format: "03/22/09 30:24 PM"

Time format: "MM/DD/YY hh:mm a" or "M/d/y HH:mm A"
but not "M/d/y HH:MM A" or "m/d/y HH:mm A"

Source format: "Time: 2009-FEBRUARY-01T22:33:44"

Time format: "Time: 'y-MMM-d'T'h:m:s" or "'Time: ' y-MMMMM-d'T'h:m:s",
but not "'Time: 'y-MM-d'T'h:m:s"

Source format: "3. quarter 2009"

Time format: "QY" or "Q'. quarter' y or "QQ/yyyy";
but not "YQ"

9.2.2.5 Change data type - single value

Changes the data type of the incoming single value to the **Number**, **Text**, or **Date** data types.

Action	Result
Conversion of Text to Number	Numerical value of the text taking into account the decimal separator. If the decimal separator is set correctly, any thousands separator is detected automatically.
Conversion of Number to Text	Text representation of the number in the internal format, or based on the language and the specified format. You can also specify a valid number of leading zeros. If nothing is specified here, the results are formatted in the numerical format.
Conversion of Text to Date	Date value of the text in the internal format, based on the specified format and, where applicable, the language. The date must be in the AD era. Date values before the common era are not supported. The time format must be specified. The time format is made up of sequences of characters, which stand for date fields, for example, year, month, day of the week, or minute, in the relevant language;

Action	Result
	separated by separators. In addition, the corresponding language must be specified. Non-editable text must be enclosed in quotation marks. When using the Q or q symbol for quarters, all other symbols except Y and y for years are ignored. Only the order of Q/q and Y/y is relevant.
Conversion of Date to Text	Text representation of the date in the internal format, or based on the language and the specified format Non-editable text must be enclosed in quotation marks. The format and language specifications are optional. If no format is specified, the data is output in the internal date format. If no language is specified, English (EN) is applied as the default language.
Conversion of Number to Date	Date value corresponding to the value of the number as milliseconds since 01/01/1970
Conversion of Date to Number	Number of milliseconds since 01/01/1970

Internal number format

If the user is logged in in English, the number format is Anglo-Saxon style with a period as the decimal separator and at least one decimal place, but without grouping characters.

Internal date format

yyyy-Q for specifying to the nearest quarter, otherwise yyyy-MM-ddThh:mm:ss. The number of digits corresponds to the accuracy of the date, and the remaining digits are omitted. This is the transfer format.

Quarterly specifications

These are indicated by a **Q** within the section of the format that is not in single quotation marks.

Prerequisites for conversion of text into quarterly date values:

1. It is expected that a source value containing a quarterly date consists of just two sequences of figures indicating the year and the quarter. Any non-numerical characters can occur before, after and between them, for example, Quarter 04/2009.
2. The pattern uses **Y** or **y** as the symbol for the year and **Q** or **q** for the quarter, for example, quarter Q/y or Q Y.

Procedure:

1. The (first) two sequences of digits are determined from the source value.
2. The section of the format that is not enclosed between single quotation marks is used to determine whether **q / Q** or **y / Y** appears first.
3. If q or Q appears first, the first sequence of digits is interpreted as the quarter and the second as the year, otherwise the reverse.

9.2.2.5.1 Parameters

The following parameters are available.

Parameter	Description
Single value	Source: Single-value operator Data type: Date, Number, or Text Specification: Mandatory
New type	New single-value data type Default value: Text Specification: Mandatory
Format (Date type)	Time format for conversion from Date type to Text type and vice versa. The following formatting symbols are available when converting date into text: Year: y or Y Quarter: Q Month: M Calendar week: w Day of the week: E or e Day of the month: d Day of the year: D Hour: H or h Minute: m Second: s AM/PM: a Time zone: z (for example, GMT) RFC time zone: Z (for example, -0900) Era: G (must always be AD) Default value: MM/dd/yyyy
Format (Date type)	The following formatting symbols are available when converting text to date: Year: y Quarter: Q Calendar week: w Week of the month: W Day of the week: E Day of the month: d

Parameter	Description
	<p>Day of the year: D</p> <p>Hour (0-23): H</p> <p>Hour (1-24): k</p> <p>Hour AM/PM (0-11): K</p> <p>Hour AM/PM (1-12): h</p> <p>Minute: m</p> <p>Second: s</p> <p>AM/PM: a</p> <p>Time zone: z</p> <p>RFC time zone: Z</p> <p>Era: G</p> <p>Default value: yyyy-MM-dd'T'HH:mm:ss;</p> <p>Permitted separators in both cases:</p> <p>Dash/minus (-), underscore (_), slash (/), period (.), colon (:), comma (,), tab character, and space.</p> <p>Specification: Mandatory</p>
Language	<p>Language if the target format is of the Date type.</p> <p>Available languages: de and en.</p> <p>Specification: Mandatory when using names of months and names of days of the week</p>
Decimal separator	<p>Separator for the decimal places, if the target format is of the Number type.</p> <p>Default value: Comma (,)</p> <p>Specification: Mandatory</p>
Format (Number source format)	<p>Number format for the conversion of the Number type to the Text type.</p> <p>You can select predefined formats or set your own format manually.</p> <p>With manual entry, the numbers before the decimal separator must have four digits ascending and then descending, for example, 1,234.321. After this, you can add text (such as the unit "hours" or km/h).</p> <p>Default value: 1234</p> <p>Permitted separators:</p> <p>Thousands separator in German: period (.)</p> <p>Thousands separator in English: comma (,)</p> <p>Decimal separator in German: comma (,)</p>

Parameter	Description
	<p>Decimal separator in English: period (.)</p>
<p>Leading zeros</p>	<p>Number of leading zeros. The maximum number of leading zeros is the number of digits before the decimal separator.</p> <p>Example:</p> <p>Format: 1,234.12 and leading zeros: 5</p> <p>Number -> Text</p> <p>10,245 -> 00010,25</p> <p>12000,4 -> 12000,4</p> <p>89,7 -> 00089,70</p>
<p>Specify type</p>	<p>Automatically specifies the data type of the source values.</p> <p>If the content of a column does not correspond to its assigned data type, a row is created in the operator which specifies the data type determined for this column Vice-versa, settings (lines) are removed from the operator, which would reset the data type of a column already typified correctly.</p>
<p>Encoding (Text source format)</p>	<p>Specifies the encoding of special characters, for the conversion of the Text type to the Text type, for example, "/", "&", "?".</p> <p>Apply UTF-8 encoding: encodes the entire text, using UTF-8 codes</p> <p>Decode UTF-8: decodes the entire text, using UTF-8 codes</p> <p>Apply URL encoding: keeps the URL specific characters</p> <p>For UTF-8 encoding/decoding the special characters must be masked in a valid URL. Only select this option if you are sure that all sections that make up the URL are already masked.</p>

The characters in the time format can be combined in any order and repeated any number of times.

Exceptions:

For a month, the number of characters must be ≥ 3 (MMM or MMMM) if the month is specified in text format (JAN, FEB, etc.) and < 3 if it is specified as a figure. In this case, a language must also be specified so that the name of the month can be transformed correctly.

For a year format such as 2009, y can be specified any number of times, that is, yy and yyyy return 2009.

For a year format such as 09, however, yyyy returns the year 9 and yy the year 2009.

When formatting date values as days of the week for a date to text conversion, an e/E number < 4 returns the day abbreviations (MON, TUE, etc.), while e/E = or > 4 returns the full name of the day.

Only the month (M), minute (m), time zone (z), RFC time zone (Z) and calendar week (w) are case-sensitive.

When converting text to date, if the Q or q symbol is used for quarters all other symbols except Y and y are ignored. Only the order of Q/q and Y/y is decisive then.

For the reverse conversion from date to text, the Q/q can be combined with any other symbols, but may only occur once (not QQ/yy)

Quarter entries are currently only possible in the form YYYY-Q. These strings can only contain the year, separator, and quarter.

All other strings must be enclosed in single quotation marks ('). Spaces can be inside or outside, for example, 'On' dd.MM.yy 'at' hh:mm, or 'On 'dd.MM.yy' at 'hh:mm' '.

The space pattern in the source and target format must match, for example, "2 .3 .09" -> "d .M .y" but not "2. 3. 09" -> "d .M .y".

9.2.2.5.2 Examples

Source format: "22.3.2009"

Time format: "d.M.y" or "DDDD.MM.YYYY",
but not "DD.MMM.YYYY "

Source format: "03/22/09 30:24 PM"

Time format: "MM/DD/YY hh:mm a" or "M/d/y HH:mm A"
but not "M/d/y HH:MM A" or "m/d/y HH:mm A"

Source format: "Time: 2009-FEBRUARY-01T22:33:44"

Time format: "Time: 'y-MMM-d'T'h:m:s" or "'Time: 'y-MMMMM-d'T'h:m:s",
but not "'Time: 'y-MM-d'T'h:m:s"

Source format: "3. quarter 2009"

Time format: "QY" or "Q'. quarter' y" or "QQ/yyyy";
but not "YQ"

9.2.2.6 Column to value

Takes a single value from a data feed column. If the column contains multiple values, the first value found is returned.

Determines the first value found from a column in the source table and returns this as a single value.

9.2.2.6.1 Parameters

The following parameters are available.

Parameter	Description
Source column	Name of the column whose values are filtered. Source: Source table Data type: Number, Text, Date; Specification: Mandatory

9.2.2.7 Combine data feeds

Merges two data feeds by comparing the values in key columns one row at a time. The key columns for the left and right table are defined in pairs. Several pairs of key columns can be specified. Both key columns must have the same data type.

One table is defined as the main table, to which all columns from the second table are added except for its key columns. The main table is linked to the upper left anchor point.

9.2.2.7.1 Parameters

The following parameters are available.

Parameter	Description
Left/right column	Name of the left or right key column. Source: Source tables Data type: Text, Date or Number For Text data type, the Case sensitive and Ignore spaces options are additionally available.

Parameter	Description
Options	
Include key values of left data feed	<p>Always transfers all key values from the left data feed (main data feed), regardless of whether there are matching rows in the right data feed. Rows with matching key values are merged. Rows in the right data feed whose key values do not occur in the left table are omitted.</p> <p>This option is selected by default.</p>
Include identical key values of both data feeds	<p>Transfers only the rows whose key values match in the two data feeds, and which therefore can be merged.</p>
Include key values of both data feeds	<p>Always transfers the key values from both data feeds, even if their key values do not occur in the other data feed. Rows with matching key values are merged.</p>
Allow multiple values	<p>Allows multiple occurrences of rows with identical key values in the right table. This can lead to a large number of result rows, as all combinations of the rows with identical key values are transferred to the results.</p> <p>This option is unselected by default.</p>

The key columns have the name they had in the left table.

Since the individual table columns are identified by name when being imported you need to ensure that the columns of the table area to be imported have unique names.

If other columns with identical names occur in both feeds, other than the key columns, `_L` or `_R` is appended to the names of these columns.

9.2.2.8 Concatenate data feeds

Adds the rows from the right-hand table after the final row of the left-hand table and merges columns of the same name and type.

For every row in the main table, a check is made as to whether there is a row in the right table that has the same values in all key column pairs. These rows are then combined into one row.

9.2.2.8.1 Parameters

The following parameters are available.

Parameter	Description
Left/right data feeds	Two data feeds to be combined. Specification: Mandatory

9.2.2.9 Concatenate texts

Combines the values of the specified columns or text fragments into one text.

Appends the values from the source columns or the source values to one another one row at a time, writes the results to the target column, and overwrites any existing values there. If the target column does not exist, it is created.

9.2.2.9.1 Parameters

The following parameters are available.

Parameter	Description
Text	Value to be linked. Source: Source table, single-value operator, input value, or a constant. Data type: Number, Text, Date; Specification: Optional
Target column	Name of the column to which the linked text is written. Source: Source table or constant. Data type: Text Default value: Result_1 Specification: Optional

9.2.2.10 Conditional replace

Changes the value in the specified column one row at a time if certain conditions are met.

Replaces existing values in the column with new values. Replacement must be linked to a condition, that is, you can specify whether all or at least one condition must be met. Several conditions can be specified and these are linked to each other with "AND".

9.2.2.10.1 Parameters

The following parameters are available.

Action	Result
Column	<p>Name of the column whose values are replaced.</p> <p>Source: Source table</p> <p>Data type: Number, Text, Date;</p> <p>Specification: Mandatory</p>
New value	<p>Value that replaces the value in the source column.</p> <p>Source: Column with values, single value from a feed (single-value operator), user input, or a constant.</p> <p>Default value: is equal to</p> <p>Specification: Mandatory</p>
Replace	<p>Values are replaced if one or all conditions is/are met.</p>
Source column	<p>Name of the column whose values are compared.</p> <p>Source: Source table</p> <p>Data type: Number, Text, Date;</p> <p>Specification: Mandatory</p>
Comparison operator	<p>Operator that compares the values from the source column with the comparison values.</p> <p>Available comparison operators depend on the data type of the source column.</p> <p>Default value: is equal to</p> <p>Specification: Mandatory</p>
Comparison values	<p>Values that are compared with the values from the source column.</p> <p>Source: Column with values, single value from a feed (single-value operator), user input, or a constant.</p> <p>Data type: Must be identical to that of the source column.</p> <p>Comparison value missing</p> <p>Condition met: If a comparison value is missing, the condition is assumed to be met.</p> <p>Condition not met: If a comparison value is missing, the condition is assumed to not be met.</p> <p>Specification: Mandatory</p>

9.2.2.10.2 Comparison operators

The following comparison operators are available.

Data type	Comparison operators
Number	<ul style="list-style-type: none"> Is equal to Is not equal to Is less than Is less than or equal to Is greater than Is greater than or equal to Is empty Is not empty
Text	<ul style="list-style-type: none"> Is equal to Is not equal to Starts with Ends with Contains Does not contain Is empty Is not empty
Date	<ul style="list-style-type: none"> Before After On Before or on On or after Is empty Is not empty

9.2.2.11 Convert text

Converts all characters in the source column one row at a time, based on the specified transformation rule. The transformation rule includes all rows in the selected source column.

9.2.2.11.1 Parameters

The following parameters are available.

Parameter	Description
Column	<p>Name of the column whose values are converted.</p> <p>Source: Source table</p> <p>Data type: Text</p> <p>Specification: Mandatory</p>
Conversion	<p>Transformation rule for conversion of column values:</p> <p>UPPER: Converts all characters into upper case, according to the rules of the specified language.</p> <p>LOWER: Converts all characters into lower case, according to the rules of the specified language.</p> <p>ONLY_LETTERS: Removes all figures (0-9) from the column values;</p> <p>ONLY_NUMBERS: Removes all letters from the column values.</p> <p>REMOVE_SPACES: Removes all spaces from the column values.</p> <p>REMOVE_LEADING_WHITESPACE: Removes leading whitespace from the column values.</p> <p>REMOVE_TRAILING_WHITESPACE: Removes trailing whitespace from the column values.</p> <p>Specification: Mandatory</p>
Target column	<p>Name of the column to which the conversation search result is written. This can be either a new column (typing a column name in the text field) or existing column (selecting a column from the drop-down menu).</p> <p>Data type: Text</p> <p>Default value: Result_1</p> <p>Specification: Optional</p> <p>If the target column is identical to the source column, the values in the source column are overwritten.</p>

9.2.2.12 Copy data feeds

Creates up to four independent copies of a data feed.

9.2.2.12.1 Parameters

The following parameters are available.

Parameter	Description
Data feed	Data feed to be copied. Specification: Mandatory

9.2.2.13 Copy single value

Creates an independent copy of a single value, without changing the input value.

9.2.2.13.1 Parameters

The following parameters are available.

Parameter	Description
Single value	Single value to be copied. Specification: Mandatory

9.2.2.14 Delete column

Deletes the specified columns from the data feed.

9.2.2.14.1 Parameters

The following parameters are available.

Action	Result
Column	Name of the column to be deleted. Source: Source table Data type: Date, Number, or Text Specification: Mandatory

9.2.2.15 Duplicate column

Copies the specified columns from the data feed to new or existing columns of the same type.

It is possible to create multiple copies of a column, but the target columns must have different names.

If the target column does not exist, it is created. If it does exist, it is replaced. Regardless of their type, all columns can be duplicated.

9.2.2.15.1 Parameter

The following parameters are available.

Action	Result
Source column	Name of the column to be duplicated. Source: Source table Data type: Date, Number, or Text Specification: Mandatory
Target column	Name of the new or existing column. Source: Constant Data type: Corresponds to source column. Specification: Mandatory

9.2.2.16 Extract text

Creates an extract from each value in a text column starting from the specified position (start index), and with the specified length, and writes the result to a target column.

Searches the source column at the specified start index, using the specified length for the string, and displays it in the target column. Start index and length must be ≥ 0 , otherwise an empty entry appears in the target column.

9.2.2.16.1 Parameters

The following parameters are available.

Parameter	Description
Column	Name of the column whose values are searched. Source: Source table Data type: Text Specification: Mandatory
Start index	Start position of the string to be extracted.

Parameter	Description
	Source: Source column, single-value operator, input value, or constant. Size >= 0; Data type: Number Specification: Mandatory
Length	Number of characters in the string to be extracted. Source: Source column, single-value operator, user input, or constant. Number of characters >= 0 Data type: Number Specification: Mandatory
Target column	Name of the column to which the search result is written. This can be either a new column (typing a column name in the text field) or existing column (selecting a column from the drop-down menu). Data type: Number Default value: Result_1 Specification: Mandatory

9.2.2.17 Filter by date

Searches a date column for the latest or earliest date and transfers these rows to the results table. All other rows are filtered out. The search can be limited to specific dimensions. If one or more dimensions are specified, the operator determines the feed row with the earliest or latest date within the feed rows with identical dimension values and transfers this to the result table. If there are several feed rows with the earliest or latest date, all of them are transferred to the results table.

9.2.2.17.1 Parameters

The following parameters are available.

Parameter	Description
Source column	Name of the source column for which the earliest or latest date values are determined. Source: Source table Data type: Date Specification: Mandatory
Earliest/latest	Determines the earliest or latest date values in the source column.

Parameter	Description
date	Default value: Earliest date
Dimension column	<p>Dimension for which the earliest or latest date values are determined. Acts as a filter to restrict the values determined.</p> <p>Data type: Text</p> <p>Specification: Mandatory</p> <p>Multiple dimension columns can be set.</p>

9.2.2.18 Filter rows

Filters the data feed one row at a time using specific conditions.

Column values of the Number, Text, or Date type are either let passed or blocked. An appropriate filter criterion can be selected depending on the data type.

The operator allows the processing of single values and value lists. You can connect the single-value user inputs and the user inputs of List type, see User input operators (page 357).

Action	Result
Action	<p>Executed if particular conditions are met.</p> <p>Possible actions:</p> <ul style="list-style-type: none"> ▪ Let values pass (from source table) ▪ Block values (from source table) <p>if</p> <ul style="list-style-type: none"> ▪ all conditions are met ▪ one condition is met <p>Default value: Let values pass if all conditions are met.</p> <p>Specification: Mandatory</p>
Column	<p>Name of the column whose values are filtered.</p> <p>Source: Source table</p> <p>Data type: Number, Text, Date;</p> <p>Specification: Mandatory</p>
Comparison operator	<p>Operator that compares the values from the source column with the comparison values.</p> <p>Available comparison operators depend on the data type of the source column.</p> <p>Default value: is equal to</p> <p>Specification: Mandatory</p>

Action	Result
Comparison values	<p>Values that are compared with the values from the source column.</p> <p>Source: Source table, single-value operator, user input, user input (List), or a constant</p> <p>Data type: Must be identical to that of the source column.</p> <p>Comparison value missing</p> <ul style="list-style-type: none"> ▪ Condition met: If a comparison value is missing, the condition is assumed to be met. ▪ Condition not met: If a comparison value is missing, the condition is assumed to not be met. <p>Specification: Mandatory</p>
All conditions should match.	<p>Combines the specified filter conditions. If the option is not enabled, the result will be a concatenation of all single filter results.</p>

9.2.2.18.1 Parameters

The following parameters are available.

Parameter	Description
Number	<p>Is equal to</p> <p>Is not equal to</p> <p>Is less than</p> <p>Is less than or equal to</p> <p>Is greater than</p> <p>Is greater than or equal to</p> <p>Is empty</p> <p>Is not empty</p>
Text	<p>Is equal to</p> <p>Is not equal to</p> <p>Starts with</p> <p>Ends with</p> <p>Contains</p> <p>Does not contain</p>

Parameter	Description
	Is empty Is not empty
Date	Before After In Before or on On or after Is empty Is not empty

COMPARISON OPERATORS PROCESSING VALUE LISTS

If a value list is processed (that is, the comparison value is a value list), either all filter values (ALL) or just a few filter values (ANY) can be applied for the filter condition (see table below). Click the  **Operator condition** icon to view the settings. The settings are also displayed in the tool tip.

- For filter conditions using the comparison operators is equal to, starts with, ends with, contains, and on the predefined applied list values are set to ANY.
- For filter conditions using the comparison operators is not equal to, does not contain, is less than, is greater than, and after the predefined applied list values are set to ALL.
- The setting of the applied filter values can be changed for all comparison operators except is equal to, is not equal to, and on. Click the  **Operator condition** icon and select the settings required.

9.2.2.19 Find text index

Finds the specified search text in a search column and writes the position of the text found to a numerical target column.

If the search text is not found, the position is -1.

9.2.2.19.1 Parameters

The following parameters operators are available.

Parameter	Description
Column	Name of the column whose values are searched. Source: Source table

Parameter	Description
	Data type: Text Specification: Mandatory
Search text	String for which the search is performed. Source: Column values from source table, single value from a feed (single-value operator), input value, or a constant. Data type: Text Specification: Mandatory
Target column	Name of the column to which the search result is written. This can be either a new column (typing a column name in the text field) or existing column (selecting a column from the drop-down menu). Data type: Number Default value: Result_1 Specification: Mandatory
First/last hit	If multiple results are found, the first or last hit is taken as the search result.

9.2.2.20 Goal accomplishment

Calculates the degree of goal accomplishment of column values one row at a time, based on the rating and the two planned values for 100% and 0%.

9.2.2.20.1 Parameters

The following parameters are available.

Parameter	Description
Calculation column value	Name of the column for which the goal accomplishment is calculated. Source: Source table Data type: Number Specification: Mandatory
Rating	Rating of the column values for which the goal accomplishment is calculated. Valid values: Positive or Negative Positive: Higher values are assessed as positive, for example, sales revenue Negative: Higher values are assessed as negative, for example, process

Parameter	Description
	throughput time. Data type: Text Default value: Positive: Specification: Mandatory
100% relates to	Target values that are compared with the source values. Source: Source table, single-value operator, input value, or a constant. Data type: Number Specification: Mandatory Goal accomplishment depends on the rating: Positive rating: Source values \geq target values Negative rating: Source values \leq target values
0% relates to	Target values that are compared with the source values. Source: Source table, single-value operator, input value, or a constant. Data type: Number Specification: Mandatory Goal accomplishment depends on the rating: Positive rating: Target values \leq source values Negative rating: Target values \geq source values
Target column	Name of the column to which the result is written. Source: Source table or constant. Default value: Result_1 Data type: Number Specification: Optional

9.2.2.21 Insert column

Inserts new columns of the **Text**, **Number**, or **Date** data type into the data feed. Each of the columns can be populated with an initial value.

9.2.2.21.1 Parameter

The following parameters are available.

Action	Result
Column name	Name of the new column.

Action	Result
	Source: Constant Data type: Date, Number, or Text Specification: Mandatory
Type	New column data type; Date, Number, or Text. Default value: Text; Specification: Mandatory
Create numeric enumeration	Fills a new column with ascending values. The values start at 1 or the value entered in the Value input box and increase by a value of 1 in each subsequent row. Specification: Optional If the option is enabled the Value box is disabled and any (default) value already entered or selected is deleted. Incoming connections for dynamic values are ignored.
Value	Initial value of the new column. Source: User input, constant or source table Data type: Depends on the data type of the source column. Specification: Optional

9.2.2.22 Merge single texts

Concatenates multiple text values.

By default, the number of characters in a text cell is limited to 2,000. This limitation applies to text cells that are part of a feed result (also of a partial result). The limitation does not apply to individual values during feed calculation.

9.2.2.22.1 Parameters

The following parameters are available.

Parameter	Description
Text	Any strings

Parameter	Description
	Source: User input, single-value operator, or constant Data type: Text Specification: Optional

Example

An SQL statement is assembled by the **Merge single texts** operator. As long as it is handled as an individual value, this value can exceed the 2,000 characters. As soon as it is used in a table, however, it is automatically shortened to 2,000 characters.

9.2.2.23 Move date

Moves a date by a specified amount of time in a given direction and writes the results to a target column.

A date can only be moved by an amount of time, the unit of which is the same as or less accurate than the unit of the date itself. If the format of the moving period is more accurate than the format of the source date, the source date is retained. If you move a date by quarters, it is moved by three months for every quarter.

If a date accurate to the nearest day with a number of days > 28 is moved to a month that has fewer days, the result is the last day of the target month.

Example

You can move a date accurate to the nearest month by months, quarters or years, but not by days. A date accurate to the nearest year can only be moved by years, a date accurate to the smallest unit by any unit.

9.2.2.23.1 Parameter

The following parameters are available.

Action	Result
Source column	Name of the source column whose date values are moved. Source: Source table Data type: Date Specification: Mandatory
Direction	Direction in which the date is moved. Valid values: Forward or Backward Data type: Text Default value: Forward

Action	Result
	Specification: Mandatory
Value	<p>Value by which the date is moved by the selected unit.</p> <p>Source: Manual entry, single value or source table</p> <p>Data type: Number</p> <p>Default value: 1</p> <p>Specification: Optional</p>
Unit	<p>Unit of time by which the date is moved.</p> <p>Data type: Text</p> <p>Default value: Second</p> <p>Specification: As source</p> <p>The information As source allows a move even if the unit of the date values is unknown at the time of creation or if it can vary.</p>
Target column	<p>Name of the target column to which the result is written.</p> <p>Data type: Date</p> <p>Default value: Result_1</p> <p>Specification: Optional</p> <p>The target column can be identical to the source column. The values in the target column are overwritten.</p>

9.2.2.23.2 Example

Source value	Value	Unit	Direction	Result
2009-12-24T16:23	10	YEAR	Forward	2019-12-24T16:23
2009-12-24	10	DAY	Forward	2010-01-03
2009-12-30	10	MONTH	Backward	2009-02-28
2009-11	1	QUARTER	Forward	2010-02
2009-11	1	DAY	Forward	2009-11
2009-01-01	3	As source	Forward	2011-01-04
2011-Q1	3	As source	Forward	2011-Q4

9.2.2.24 RAQL Inline

Transforms up to four data feeds into one output by applying a RAQL statement.

9.2.2.24.1 Parameters

The following parameters are available.

Parameter	Description
Data feeds	Up to four data feeds to be transformed. In the RAQL statement, the feeds are referred to as feed1, feed2, feed3, and feed4. At least one feed must be connected.
RAQL query	Input field to enter any RAQL statement.
Insert parameter ()	Inserts user defined input parameters at cursor position. The button is only clickable if at least one user input parameter, for example, Text user input , has already been inserted in the feed definition. See RAQL and SQL statement parameters (Seite 316) for details.
Expand ()/ Collapse ()	Temporarily expand the edit box. When expanded, the box will shrink again upon clicking  , or when the editor goes out of focus.

The RAQL syntax reference can be found at RAQL Query Syntax Reference (page 367).

Built-in RAQL functions are described in Built-In RAQL Functions.

For instructions on how to create user-defined functions, see Configure, Compile, Deploy and Test User-Defined Functions.

9.2.2.24.2 Example

```
SELECT feed1.productName, feed2.supplierName FROM feed1, feed2 WHERE price <=
:maxPrice AND feed1.supplierId = feed2.id
```

9.2.2.25 Rename column

Changes the names of the specified columns from the data feed. The data type of the column is retained.

9.2.2.25.1 Parameter

The following parameters are available.

Action	Result
Column	Name of the column to be deleted. Source: Source table Data type: Date, Number, or Text Specification: Mandatory
New name	New name of the renamed column. Source: Constant Data type: Corresponds to source column. Specification: Mandatory

9.2.2.26 Replace text

Replaces text in a search column with the specified Find or Replace text one row at a time, or writes the text to a target column.

If the search text cannot be found, the search text itself is written to the target column.

9.2.2.26.1 Parameters

The following parameters are available.

Parameter	Description
Column	Name of the column whose values are searched. Source: Source table Data type: Text Specification: Mandatory
Search text	String for which the search is performed. Source: Column values from source table, single value from a feed (single-value operator), input value, or a constant. Data type: Text Specification: Mandatory
Replacement text	String that replaces the search text. Source: Source table, single-value operator, input value, or constant. Data type: Text Specification: Optional If no replacement text is specified, the search text found is replaced with an empty text.

Parameter	Description
Target column	Name of the column to which the search result is written. This can be either a new column (typing a column name in the text field) or existing column (selecting a column from the drop-down menu). Data type: Number Default value: Result_1 Specification: Mandatory
First/last hit	IF multiple results are found, the first, last, or all hits is/are replaced. The specification relates to occurrence within the individual rows of the search column and not to the sequence of rows, that is, NOT "First row", "Last row" and "All rows".

9.2.2.27 Round up/down

Rounds the values from a numerical source column to the specified number of decimal places (accuracy), writes the results to the target column, and overwrites any existing values there. If the target column does not exist, it is created.

If the accuracy itself is specified as a decimal number, the decimal places are ignored, that is, the integer value is used. Values that already have the same number or fewer decimal places than specified remain unchanged.

When rounding, the value is rounded down if the next decimal place is < 5, otherwise it is rounded up.

9.2.2.27.1 Parameter

The following parameters are available.

Action	Result
Source column	Name of the source column whose values are rounded. Source: Source table Data type: Number Specification: Mandatory
Precision	Numerical value specifying the number of decimal places; Source: Source table, single-value operator, input value, or a constant. Data type: Number Specification: Mandatory
Target column	Name of the column to which the result is written. The column name can be

Action	Result
	<p>transferred from the source table or freely entered.</p> <p>Data type: Number</p> <p>Default value: Result_1</p> <p>Specification: Optional</p>

9.2.2.28 Round up/down date

Converts date values from a date column to a rougher time unit and writes the results to a target column.

9.2.2.28.1 Parameter

The following parameters are available.

Action	Result
Source column	<p>Name of the source column whose values are rounded.</p> <p>Source: Source table</p> <p>Data type: Date</p> <p>Specification: Mandatory</p>
Precision	<p>Accuracy of the new date format, defined by the unit: Year, Quarter, Month, Day, Hour, Minute, or Second, and Interval: Depending on the selected unit, for example, 5 minutes or 1 year</p> <p>Data types: Numeric, Text</p> <p>Default values: 1, Minute</p> <p>If the accuracy of the source column is less accurate or the same as the target column format, the original value is retained.</p> <p>The date values are rounded according to the selected interval. Only the unit to be rounded is taken into account, for example, when rounding to minutes, the seconds are ignored.</p> <p>Rounding type: Specifies how the selected time interval is to be rounded.</p> <p>Round up for half an interval: Automatically rounds up above an interval value higher than or equal to half of the interval value</p> <p>Round down for half an interval: Automatically rounds down below an interval value lower than or equal to half of the interval value</p> <p>Always round up: Always rounds up, regardless of the interval value</p> <p>Always round down: Always rounds down, regardless of the interval value</p>

Action	Result
Target column	<p>Name of the target column to which the converted date is written</p> <p>Data type: Date</p> <p>Default value: Result_1</p> <p>Specification: Optional</p> <p>The target column can be identical to the source column. The values in the target column are overwritten.</p> <p>If the target column is not of the Date type, it is replaced by a new date column.</p>

9.2.2.28.2 Examples

Source value	Accuracy	Result
2009-12-24T16:23	Day	2009-12-24
2009-12-24T16:23	Hour	2009-12-24T16
2009-12-24	Month	2009-12
2009-12-24	Quarter	2009-Q4
2009-12-24	Year	2009

Rounding

2010-08-06T17:15:27	10 seconds	2010-08-06T17:15:30
2010-08-06T17:07:00	15 minutes	2010-08-06T17:00
2010-08-06T17:18:00	15 minutes	2010-08-06T17:15
2010-08-06T02:18:04	4 hours	2010-08-06T04:00:00

Round up

Source value	Accuracy	Result
2010-02-28T23:07:00:00 AM	15 minutes	2010-02-28T23:00
2010-02-28T23:07:30	15 minutes	2010-02-28T11:15:00 PM
2010-02-28T23:30:00	1 hour	2010-03-01T00

Round down

Source value	Accuracy	Result
1970:01:01T09:00:01	6 hours	1970:01:01T12:00:00

Source value	Accuracy	Result
1970:01:01T09:00:00	6 hours	1970:01:01T 06 :00:00

Always round up

Source value	Accuracy	Result
2010-02-28T 11:15:00 PM	15 minutes	2010-02-28T 11:15:00 PM
2010-02-28T 11:15:01 PM	15 minutes	2010-02-28T 11:30:00 PM
2010- 02-28T20:00:01	6 hours	2010- 03-01T00

Always round down

Source value	Accuracy	Result
2010-02-28T 11:15:00 PM	15 minutes	2010-02-28T 11:15:00 PM
2010-02-28T 11:14:59 PM	15 minutes	2010-02-28T 23:00
2010- 03-01T05:59:59	6 hours	2010- 03-01T00

9.2.2.29 Runtime info

Provides system information on the logged-in user or the current date. The operator can also generate a random number.

The value type of the resulting single value changes accordingly.

9.2.2.29.1 Parameters

The following parameters are available.

Parameter	Description
Information type	<p>The single-value operator can return the following types of information:</p> <ul style="list-style-type: none"> User data Today's date Random number <p>Default value: User data</p> <p>Specification: Mandatory</p>
Property	<p>Properties of the logged-in user; displayed if User data is selected as the information type.</p>

Parameter	Description
	<p>The following values can be selected: Login, First name, Last name, E-mail and Language</p> <p>Default value: User name</p> <p>Specification: Mandatory</p>
Precision	<p>Specifies the accuracy of the date, displayed if Today's date is selected as the information type.</p> <p>The following values can be selected: Minute, Hour, Day, Month, Year</p> <p>Default value: Day</p> <p>Specification: Mandatory</p>
Number range	<p>Number range of the random number; displayed if Random number is selected as the information type.</p> <p>The following values can be selected: Integers, Floating point numbers</p> <p>Default value: Integers</p> <p>Specification: Mandatory</p>
Upper/lower limit	<p>Upper or lower limit of the value range for the random number; displayed if Random number is selected as the information type.</p> <p>Default value: 0 and 10</p> <p>Specification: Mandatory</p>

9.2.2.30 Value to column

Converts an individual value into a column so that it can be connected to an operator.

Creates a feed table from a single-value operator with a column of the source operator type and a row containing the value of the source operator.

9.2.2.30.1 Parameters

The following parameters are available.

Parameter	Description
Target column	<p>Name of the column to which the conversion result is written.</p> <p>The name of the column is initially Result_1. You can change the name manually if required.</p> <p>Data type: Text, Date or Number</p> <p>Default value: Result_1</p>

Parameter	Description
	Specification: Optional

9.2.3 User input operators

A user input operator is an interface to a data feed, allowing a user to enter data manually in a dashboard. The following user input operators are available in the data feed editor.

9.2.3.1 Date user input

The date user input enables the dynamic entry of date values in data feed processing. User input is an interface to a data feed, allowing a user to enter data manually in a dashboard.

The input has the format yyyy-MM-dd'T'HH:mm:ss (up to the required accuracy) or yyyy-'Q'Q.

9.2.3.1.1 Parameters

The following parameters are available.

Parameter	Description
Name	Name of user input Source: Constant Data type: Text Specification: Optional The names of the individual user input must be unique within the feed definition.
Debug value	Value used for a test calculation in the Feed Editor. Source: Constant Data type: Date Specification: Optional
Preview value	The value is used if the user does not provide any input. Source: Constant Data type: Date Specification: Optional

9.2.3.2 Date user input (List)

The date user input enables the dynamic entry of date values in data feed processing. The operator is an interface to a data feed that can process multiple values (lists of values) at the same time. It enables the multiple selection in widgets.

The user input (list) operator can only be connected to operators that support the processing of list values (multiple selection). These operators are:

- Filter rows
- Data feed
- PPM
- JDBC
- Terracotta DB

The date user input has the format yyyy-MM-dd'T'HH:mm:ss (up to the required accuracy) or yyyy-'Q'Q.

9.2.3.2.1 Parameters

The following parameters are available.

Parameter	Description
Name	Name of user input Source: Constant Data type: Text Specification: Optional The names of the individual user input must be unique within the feed definition.
Debug value	Value used for a test calculation in the Feed Editor. Source: Constant Data type: Date Specification: Optional
Preview value	The value is used if the user does not provide any input. Source: Constant Data type: Date Specification: Optional
Edit	Enables you to enter the relevant debug and preview values in a table.

9.2.3.3 Number user input

The number user input enables the dynamic entry of numerical values in data feed processing. User input is an interface to a data feed, allowing a user to enter data manually in a dashboard. The input is done with a period (.) as the decimal separator and with no thousand grouping character (for example, 1234.56).

9.2.3.3.1 Parameters

The following parameters are available.

Parameter	Description
Name	Name of user input Source: Constant Data type: Text Specification: Optional The names of the individual user input must be unique within the feed definition.
Debug value	The values is used for a test calculation in the Feed Editor. Source: Constant Data type: Number Specification: Optional
Preview value	The value is used if the user does not provide any input. Source: Constant Data type: Number Specification: Optional

9.2.3.4 Number user input (List)

The number user input enables the dynamic entry of numerical values in data feed processing. User input is an interface to a data feed, allowing a user to enter data manually in a dashboard. The input is done with a period (.) as the decimal separator and with no thousand grouping character (for example, 1234.56).

The user input (list) operator can only be connected to operators that support the processing of list values (multiple selection). These operators are:

- Filter rows
- Data feed
- JDBC

- Terracotta DB

9.2.3.4.1 Parameters

The following parameters are available.

Parameter	Description
Name	Name of user input Source: Constant Data type: Text Specification: Optional The names of the individual user input must be unique within the feed definition.
Debug value	List of value used for a test calculation in the feed editor. Source: Constant Data type: Number Specification: Optional
Preview value	List of values used if the user does not provide any input. Source: Constant Data type: Number Specification: Optional
Edit	Enables you to enter the relevant debug and preview values in a table.

9.2.3.5 Text user input

The text user input enables the dynamic entry of text in data feed processing. User input is an interface to a data feed, allowing a user to enter data manually in a dashboard.

9.2.3.5.1 Parameters

The following parameters are available.

Parameter	Description
Name	Name of user input Source: Constant Data type: Text Specification: Optional

Parameter	Description
	The names of the individual user input must be unique within the feed definition.
Debug value	Value used for a test calculation in the Feed Editor. Source: Constant Data type: Text Specification: Optional
Preview value	The value is used if the user does not provide any input. Source: Constant Data type: Text Specification: Optional

9.2.3.6 Text user input (List)

The text user input enables the dynamic entry of text in data feed processing. User input is an interface to a data feed, allowing a user to enter data manually in a dashboard.

The user input (list) operator can only be connected to operators that support the processing of list values (multiple selection). These operators are:

- Filter rows
- Data feed
- PPM
- JDBC
- Terracotta DB

9.2.3.6.1 Parameters

The following parameters are available.

Parameter	Description
Name	Name of user input Source: Constant Data type: Text Specification: Optional The names of the individual user input must be unique within the feed definition.
Debug value	Value used for a test calculation in the Feed Editor. Source: Constant Data type: Text

Parameter	Description
	Specification: Optional
Preview value	The value is used if the user does not provide any input. Source: Constant Data type: Text Specification: Optional
Edit	Enables you to enter the relevant debug and preview values in a table.

9.3 Analytics with RAQL

9.3.1 About the Real-Time Analytical Query Language

MashZone NextGen provides a textual query language called Real-Time Analytical Query Language (RAQL). With RAQL you can conduct analytics to derive insights from data coming from multiple data sources.

9.3.1.1 Scope of RAQL

RAQL is based on SQL, which is the common standard for querying structured data. RAQL provides the common subset of SQL functionality for reading and exploring data, including for example filtering, joining, and aggregation. It also includes extensions in the form of additional analytics functions so that users have a rich set of options for exploring and analyzing data. RAQL offers as well a straight-forward way to provide your own user-defined functions to meet new or unique analysis needs.

9.3.1.2 RAQL in MashZone NextGen

RAQL is available as query interface in the following building blocks of MashZone NextGen:

In feed editor RAQL inline operator is available as part of the data transformation operators. Using that intermediate operator you can connect to one or more data sources and run a RAQL query against them. See also RAQL Inline (Seite 350).

In feed editor and dashboard editor you can select the BigMemory operator and run a RAQL query against data available in a BigMemory server. See also BigMemory (page 297).

In feed editor and dashboard editor you can select the Terracotta DB operator and run a RAQL query against data available in a Terracotta DB server. See also Terracotta DB (page 311).

RAQL queries are processed within a dedicated analytical engine, which is an integral part of the MashZone NextGen server.

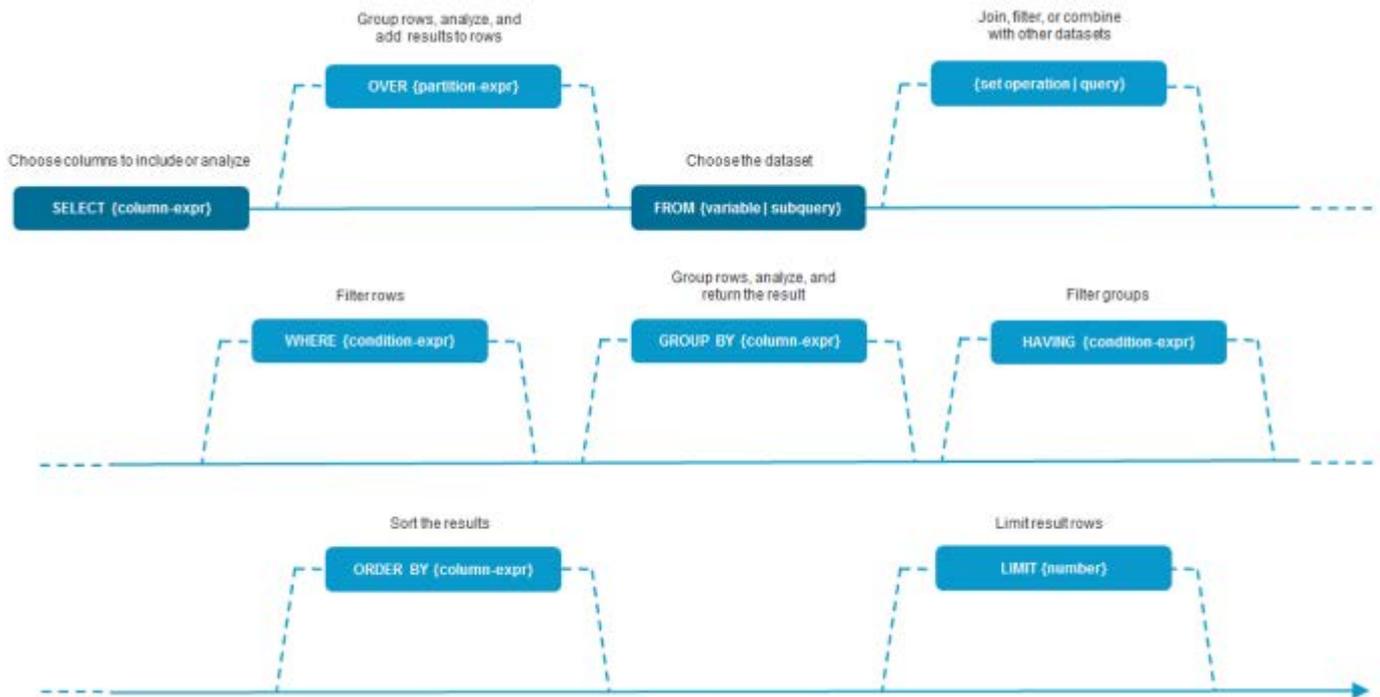
9.3.2 Working with RAQL queries

RAQL is an easy-to-learn query language for querying and analyzing multiple data sources.

9.3.2.1 Query Syntax of RAQL

Once you have selected a data source, you can use a RAQL query to analyze the data. RAQL is based on SQL and additionally equipped with functional extensions which are not defined in SQL. The subsequent documentation focuses on how to work with RAQL and with those functional extensions. For a detailed introduction to SQL we refer to the standard SQL text books.

RAQL provides the typical SQL clauses for reading and exploring data. Some SQL clauses and syntax are not found in RAQL. A RAQL query is composed of clauses:



See Language Reference (page 365) for a synopsis of the valid expressions for each of these clauses.

The most basic form of a RAQL query includes the SELECT and the FROM clause.

The SELECT clause determines which columns to include in the result and can also perform analysis when it is used with either the OVER clause or the GROUP BY clause. The FROM clause determines which dataset to query, or can define a subquery to use as the source of data.

Set operation clauses allow the query to retrieve an additional dataset, using another query, and then join, combine or filter these datasets to derive a more complex dataset. The JOIN set operation, for example, matches rows in both datasets based on a condition and adds columns from both datasets to the joined row.

The WHERE clause filters rows from the dataset. ORDER BY sorts the result rows. LIMIT determines the maximum number of result rows that the query can return.

The OVER and GROUP BY clauses both group dataset rows into different sets based on an expression. These groups determine the scope of rows that are used in analytic functions in the SELECT clause. The HAVING clause filters the set of groups that are returned in a GROUP BY clause.

OVER and GROUP BY are mutually exclusive as they have different effects on the data returned by the query. The OVER clause performs calculations and adds the calculations as additional columns to each row. GROUP BY instead performs calculations and returns just the calculations for each group.

Most RAQL clauses also support the use of functions within their expressions. RAQL functions come in two varieties:

- Plain functions, that perform some simple transformation to the values of a column for each row, such as UPPER to change text to upper-case.
- Analytic functions, more commonly known as aggregate or window functions, perform calculations using multiple rows in a group, partition or window defined in the OVER or GROUP BY clauses.
- Aggregate analytic functions use all rows in the current scope, such as SUM, while window analytic functions use specific rows, such as ROW_NUMBER. These functions include simple arithmetic as well as statistical functions or other analysis algorithms.

RAQL provides a set of function extensions (plain and analytical) as well as a way for you to define your own functions. See RAQL extensions (page 404) for more information.

9.3.2.2 Data types in RAQL

MashZone NextGen focuses in dashboard development on three data types: string, date, and floating number. In addition to these basic data types, RAQL supports other simple data types when working with input data or casting data. The data types of the RAQL query results, however, are always mapped to the three basic data types, i.e. string, date, and floating number, during follow-up processing of the results. Note that the mapping of a numeric data type, e.g., BigDecimal, to a floating number may imply a loss of precision.

The following table shows the data types RAQL supports, their Java equivalent, as well as their mapping to one of the three basic MashZone NextGen types.

RAQL Datatype	Java Equivalent	Query Result Mapping
bigdecimal	java.math.BigDecimal	number
boolean	java.lang.boolean	string
byte	java.lang.byte	string
character	java.lang.character	string
datetime	java.util.Date	date
double	java.lang.Double	number

RAQL Datatype	Java Equivalent	Query Result Mapping
float	java.lang.Float	number
integer	java.lang.Integer	number
long	java.lang.Long	number
short	java.lang.Short	number
string	java.lang.String	string

9.3.2.3 Creating dynamic RAQL queries

You can create dynamic RAQL queries which use input parameters to provide the values used in the different RAQL clauses, e.g., in the WHERE clause. See also RAQL parameters (page 316) for more details.

9.3.3 Language reference

9.3.3.1 Reserved keywords

The following SQL keywords are reserved words for RAQL queries. Keywords are shown in lower case. Matching is not case sensitive.

Column and function names in queries that match reserved keywords must be delimited with double quote marks (") to prevent errors in query execution.

Note: Single quote marks (') are used to delimit literal values that are string, and thus are not valid to delimit names.

Reserved keywords

abs	false	natural	then
all	first	not	to
and	first_value	now	trailing
any	float	nth_value	trim
array	floor	ntile	true
as	following	null	trunc
asc	for	nulls	truncate
avg	from	on	unbounded
between	full	or	union

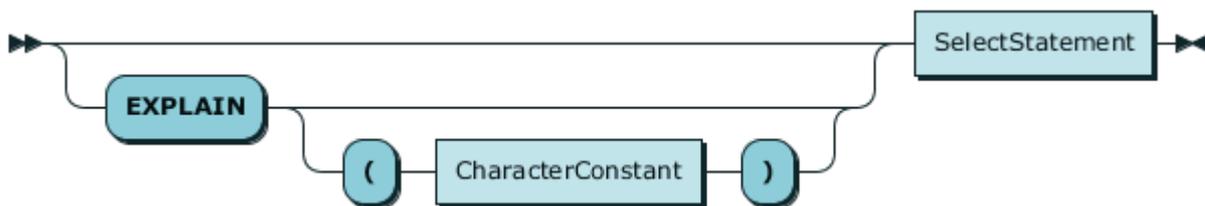
big_decimal	group	order	unknown
blob	grouping	outer	upper
boolean	having	over	using
by	hour	partition	variance
byte	hours	pattern	var_pop
case	in	power	var_samp
cast	inner	preceding	week
ceil	integer	range	weeks
ceiling	intersect	rank	when
character	is	recursive	where
character_length	join	ref	window
char_length	lag	regr_intercept	with
clob	last	regr_slope	within
corr	last_value	relative	year
count	lead	right	years
covar	leading	round	
covar_pop	left	row	
covar_samp	like	rows	
cross	limit	row_number	
cube	In	seconds	
cume_dist	long	seconds	
current	lower	select	
date	matching	sets	
day	max	shift	
days	measures	short	
default	millisecond	slide	
define	milliseconds	sql_date	
dense_rank	min	sql_time	
desc	minus	sql_timestamp	
distinct	minute	sqrts	

do	minutes	stddev	
double	mod	stddev_pop	
duration	month	stddev_samp	
else	months	string	
end		struct	
except		substring	
exp		sum	
explain			

9.3.3.2 Syntax Reference

This topic contains syntax definitions and corresponding syntax diagrams for the Real-Time Analytics Query Language (RAQL) in MashZone NextGen.

9.3.3.2.1 ExplainOrSelectStatement



Definition:

**ExplainOrSelectStatement ::= ('EXPLAIN' ('(' CharacterConstant ')')?)?
SelectStatement**

Used In:

This is the starting definition of RAQL queries, which has no parent references.

9.3.3.2.2 CharacterConstant



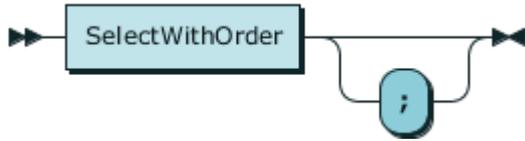
Definition:

CharacterConstant ::= S_CHAR_LITERAL | S_PARAMETER

Used In:

- ExplainOrSelectStatement (page 367)

9.3.3.2.3 SelectStatement



Definition:

SelectStatement ::= SelectWithOrder ';'?

Used In:

- ExplainOrSelectStatement (page 367)

9.3.3.2.4 SelectWithOrder



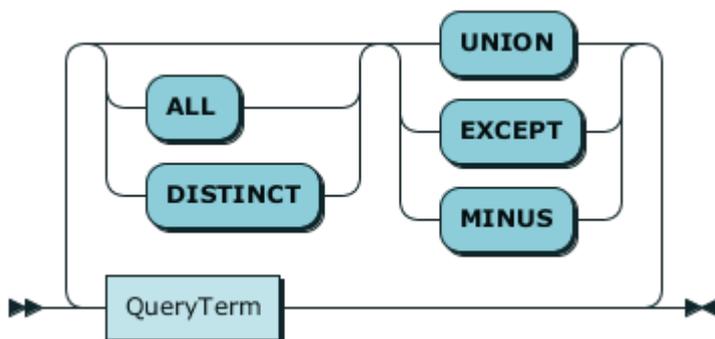
Definition:

SelectWithOrder ::= QueryExpressionBody OrderByClause? LimitClause?

Used In:

- SelectStatement (page 368)
- SourceReferenceOrSubquery (page 391)

9.3.3.2.5 QueryExpressionBody



Definition:

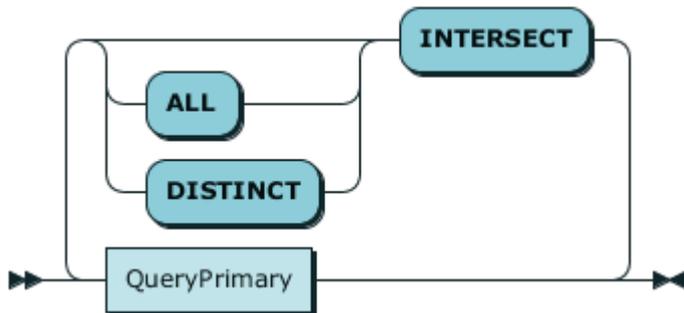
QueryExpressionBody ::= QueryTerm (('UNION' | 'EXCEPT' | 'MINUS') ('ALL' | 'DISTINCT')? QueryTerm)*

Used In:

- QueryPrimary (page 369)

- [SelectWithOrder](#) (page 368)
- [SqlRelationalOperatorExpression](#) (page 376)

9.3.3.2.6 QueryTerm



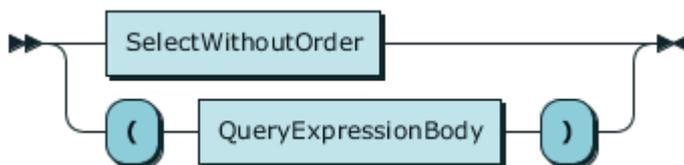
Definition:

QueryTerm ::= QueryPrimary ('INTERSECT' ('ALL' | 'DISTINCT')? QueryPrimary)*

Used In:

- [QueryExpressionBody](#) (page 368)

9.3.3.2.7 QueryPrimary



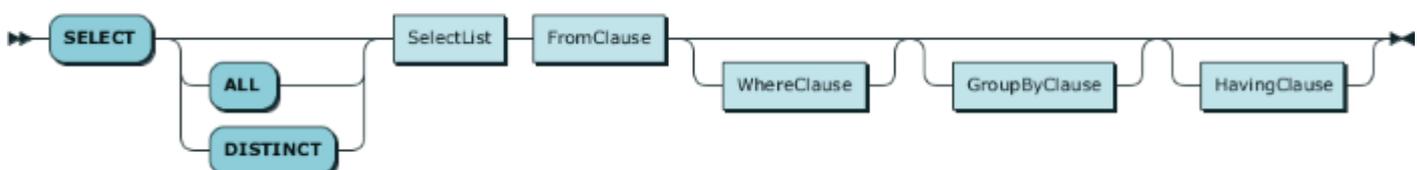
Definition:

QueryPrimary ::= SelectWithoutOrder | '(' QueryExpressionBody ')'

Used In:

- [QueryTerm](#) (page 369)

9.3.3.2.8 SelectWithoutOrder



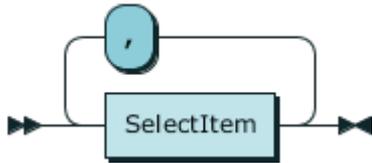
Definition:

SelectWithoutOrder ::= 'SELECT' ('ALL' | 'DISTINCT')? SelectList FromClause WhereClause? GroupByClause? HavingClause?

Used In:

- QueryPrimary (page 369)

9.3.3.2.9 SelectList



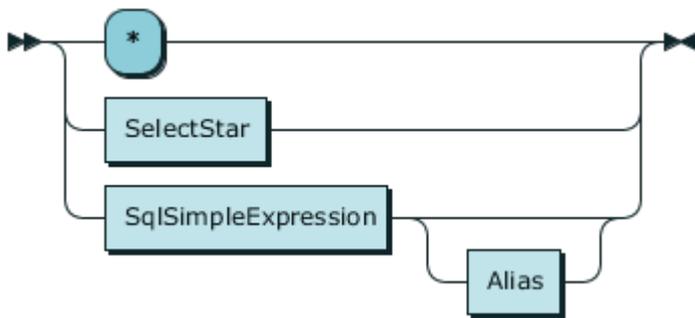
Definition:

SelectList ::= SelectItem (',' SelectItem)*

Used In:

- SelectWithoutOrder (page 369)

9.3.3.2.10 SelectItem



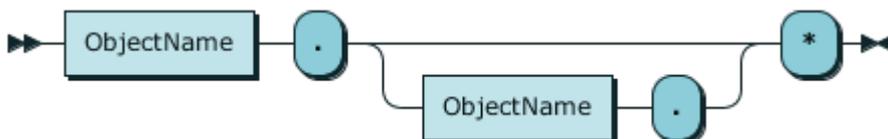
Definition:

SelectItem ::= '*' | SelectStar | SqlSimpleExpression Alias?

Used In:

- SelectList (page 370)

9.3.3.2.11 SelectStar



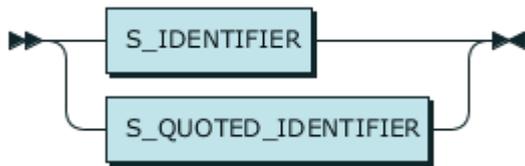
Definition:

SelectStar ::= ObjectName '.' (ObjectName '.')? '*'

Used In:

- SelectItem (page 370)

9.3.3.2.12 ObjectName



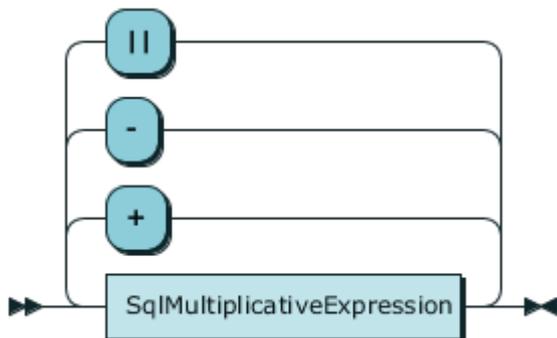
Definition:

ObjectName ::= S_IDENTIFIER | S_QUOTED_IDENTIFIER

Used In:

- Alias (page 391)
- GenericFunctionCall (page 385)
- SelectStar (page 370)
- SourceField (page 385)
- SourceReference (page 392)
- SqlPrimaryExpression (page 373)

9.3.3.2.13 SqlSimpleExpression



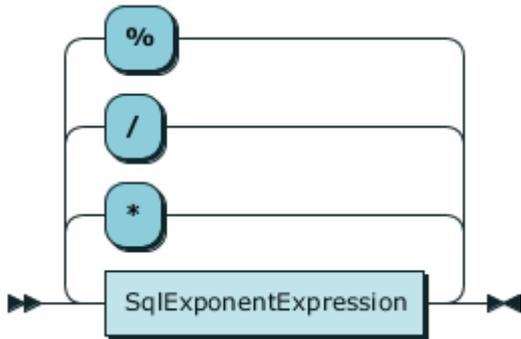
Definition:

**SqlSimpleExpression ::= SqlMultiplicativeExpression (('+' | '-' | '||')
SqlMultiplicativeExpression)***

Used In:

- BooleanCaseCall (page 382)
- SelectItem (page 370)
- SimpleCaseCall (page 382)
- SqlBetweenClause (page 377)
- SqlLikeClause (page 377)
- SqlRelationalExpression (page 375)
- SqlRelationalOperatorExpression (page 376)

9.3.3.2.14 SqlMultiplicativeExpression



Definition:

SqlMultiplicativeExpression ::= SqlExponentExpression (('*' | '/' | '%') SqlExponentExpression)*

Used In:

- SqlSimpleExpression (page 371)

9.3.3.2.15 SqlExponentExpression



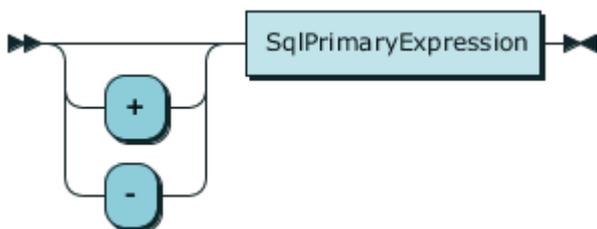
Definition:

SqlExponentExpression ::= SqlUnaryExpression ('' SqlUnaryExpression)***

Used In:

- SqlMultiplicativeExpression (page 372)

9.3.3.2.16 SqlUnaryExpression



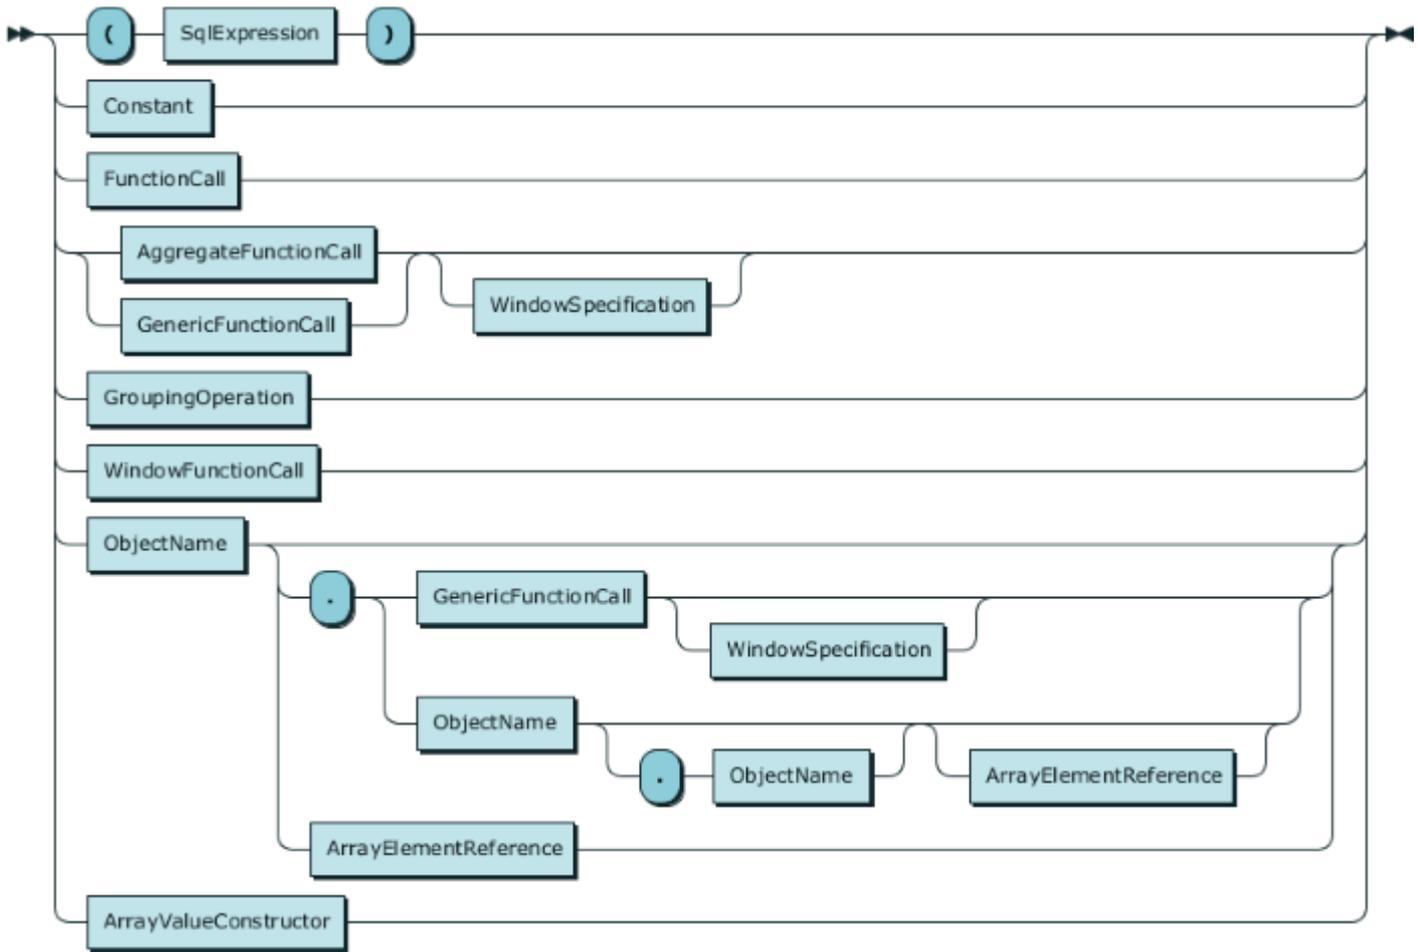
Definition:

SqlUnaryExpression ::= ('+' | '-')? SqlPrimaryExpression

Used In:

- SqlExponentExpression (page 372)

9.3.3.2.17 SqlPrimaryExpression



Definition:

```
SqlPrimaryExpression ::= '(' SqlExpression ')' | Constant | FunctionCall | (
AggregateFunctionCall | GenericFunctionCall ) WindowSpecification? |
GroupingOperation | WindowFunctionCall | ObjectName ( '.' ( GenericFunctionCall
WindowSpecification? | ObjectName ( '.' ObjectName )? ArrayElementReference? ) |
ArrayElementReference )? | ArrayValueConstructor
```

Used In:

- SqlUnaryExpression (page 372)

9.3.3.2.18 SqlExpression



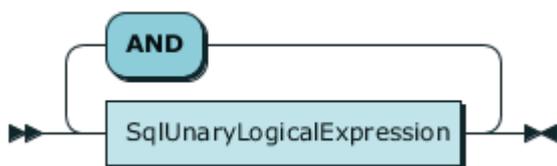
Definition:

```
SqlExpression ::= SqlAndExpression ( 'OR' SqlAndExpression )*
```

Used In:

- [AggregateFunctionCall](#) (page 384)
- [BooleanCaseCall](#) (page 382)
- [FunctionCall](#) (page 379)
- [GroupingExpressionReference](#) (page 395)
- [HavingClause](#) (page 397)
- [JoinedTable](#) (page 392)
- [OrderByClause](#) (page 386)
- [SqlExpressionList](#) (page 375)
- [SqlPrimaryExpression](#) (page 373)
- [WhereClause](#) (page 393)
- [WindowFunctionCall](#) (page 389)

9.3.3.2.19 [SqlAndExpression](#)



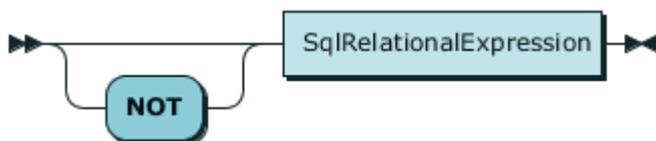
Definition:

SqlAndExpression ::= SqlUnaryLogicalExpression ('AND' SqlUnaryLogicalExpression)*

Used In:

- [SqlExpression](#) (page 373)

9.3.3.2.20 [SqlUnaryLogicalExpression](#)



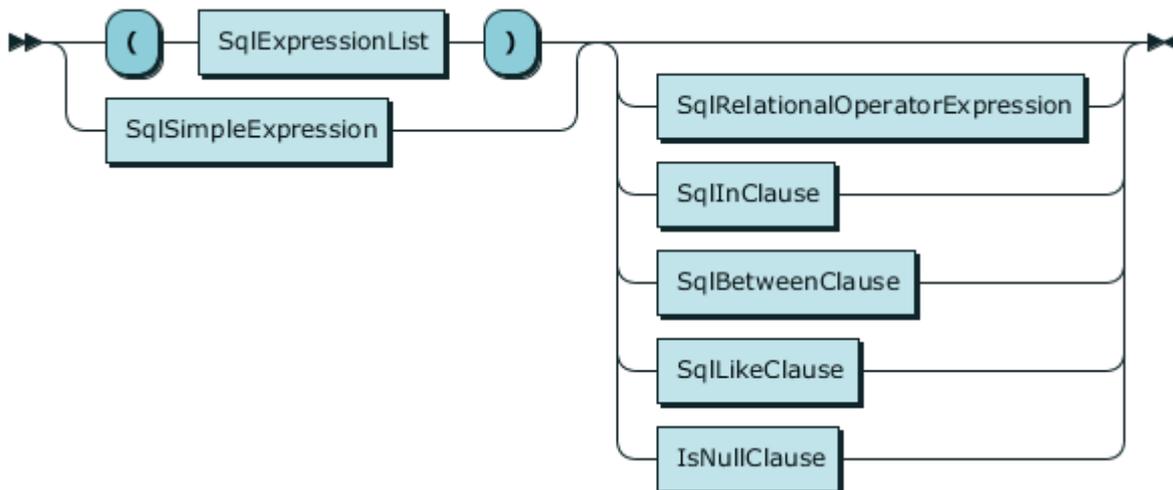
Definition:

SqlUnaryLogicalExpression ::= 'NOT'? SqlRelationalExpression

Used In:

- [SqlAndExpression](#) (page 374)

9.3.3.2.21 SqlRelationalExpression



Definition:

```
SqlRelationalExpression ::= ( '(' SqlExpressionList ')' | SqlSimpleExpression ) (
SqlRelationalOperatorExpression | SqlInClause | SqlBetweenClause | SqlLikeClause |
IsNullClause )?
```

Used In:

- SqlUnaryLogicalExpression (page 374)

9.3.3.2.22 SqlExpressionList



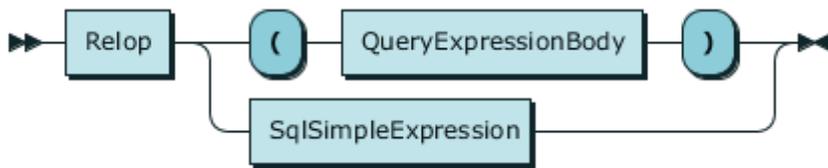
Definition:

```
SqlExpressionList ::= SqlExpression ( ',' SqlExpression )*
```

Used In:

- ArrayValueConstructor (page 390)
- GenericFunctionCall (page 385)
- GroupingOperation (page 388)
- SqlInClause (page 376)
- SqlRelationalExpression (page 375)
- WindowSpecification (page 386)

9.3.3.2.23 SqlRelationalOperatorExpression



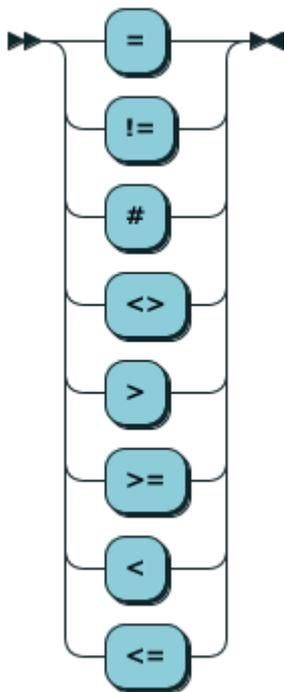
Definition:

SqlRelationalOperatorExpression ::= Relop ('(' QueryExpressionBody ')' | SqlSimpleExpression)

Used In:

- SqlRelationalExpression (page 375)

9.3.3.2.24 Relop



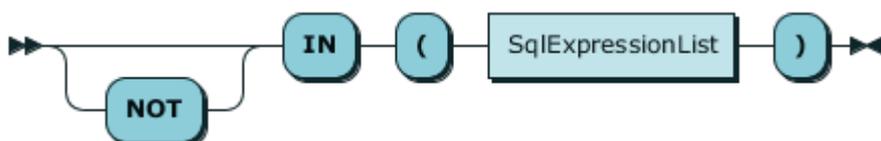
Definition:

Relop ::= '=' | '!=' | '#' | '<>' | '>' | '>=' | '<' | '<='

Used In:

- SqlRelationalOperatorExpression (page 376)

9.3.3.2.25 SqlInClause



Definition:

SqlInClause ::= 'NOT'? 'IN' '(' SqlExpressionList ')'

Used In:

- SqlRelationalExpression (page 375)

9.3.3.2.26 SqlBetweenClause



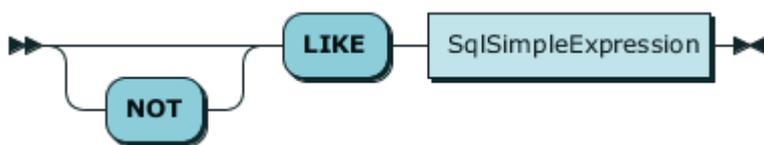
Definition:

SqlBetweenClause ::= 'NOT'? 'BETWEEN' SqlSimpleExpression 'AND' SqlSimpleExpression

Used In:

- SqlRelationalExpression (page 375)

9.3.3.2.27 SqlLikeClause



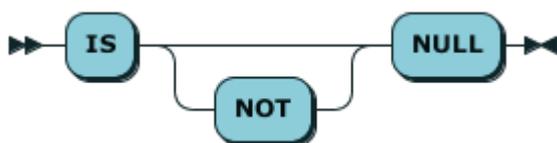
Definition:

SqlLikeClause ::= 'NOT'? 'LIKE' SqlSimpleExpression

Used In:

- SqlRelationalExpression (page 375)

9.3.3.2.28 IsNullClause



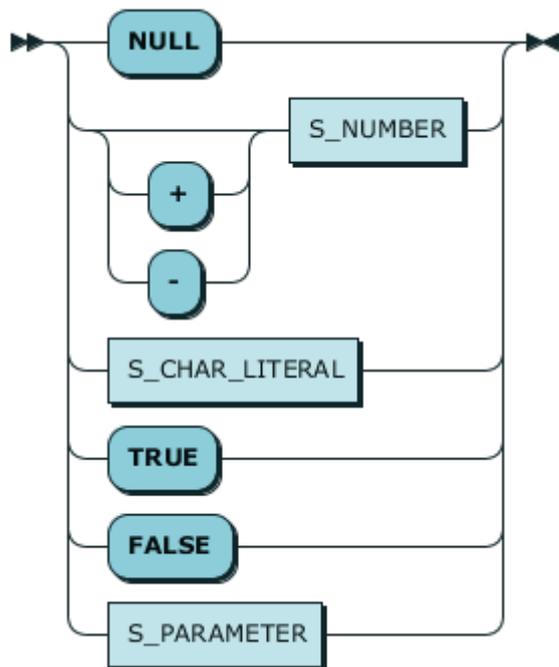
Definition:

IsNullClause ::= 'IS' 'NOT'? 'NULL'

Used In:

- SqlRelationalExpression (page 375)

9.3.3.2.29 Constant



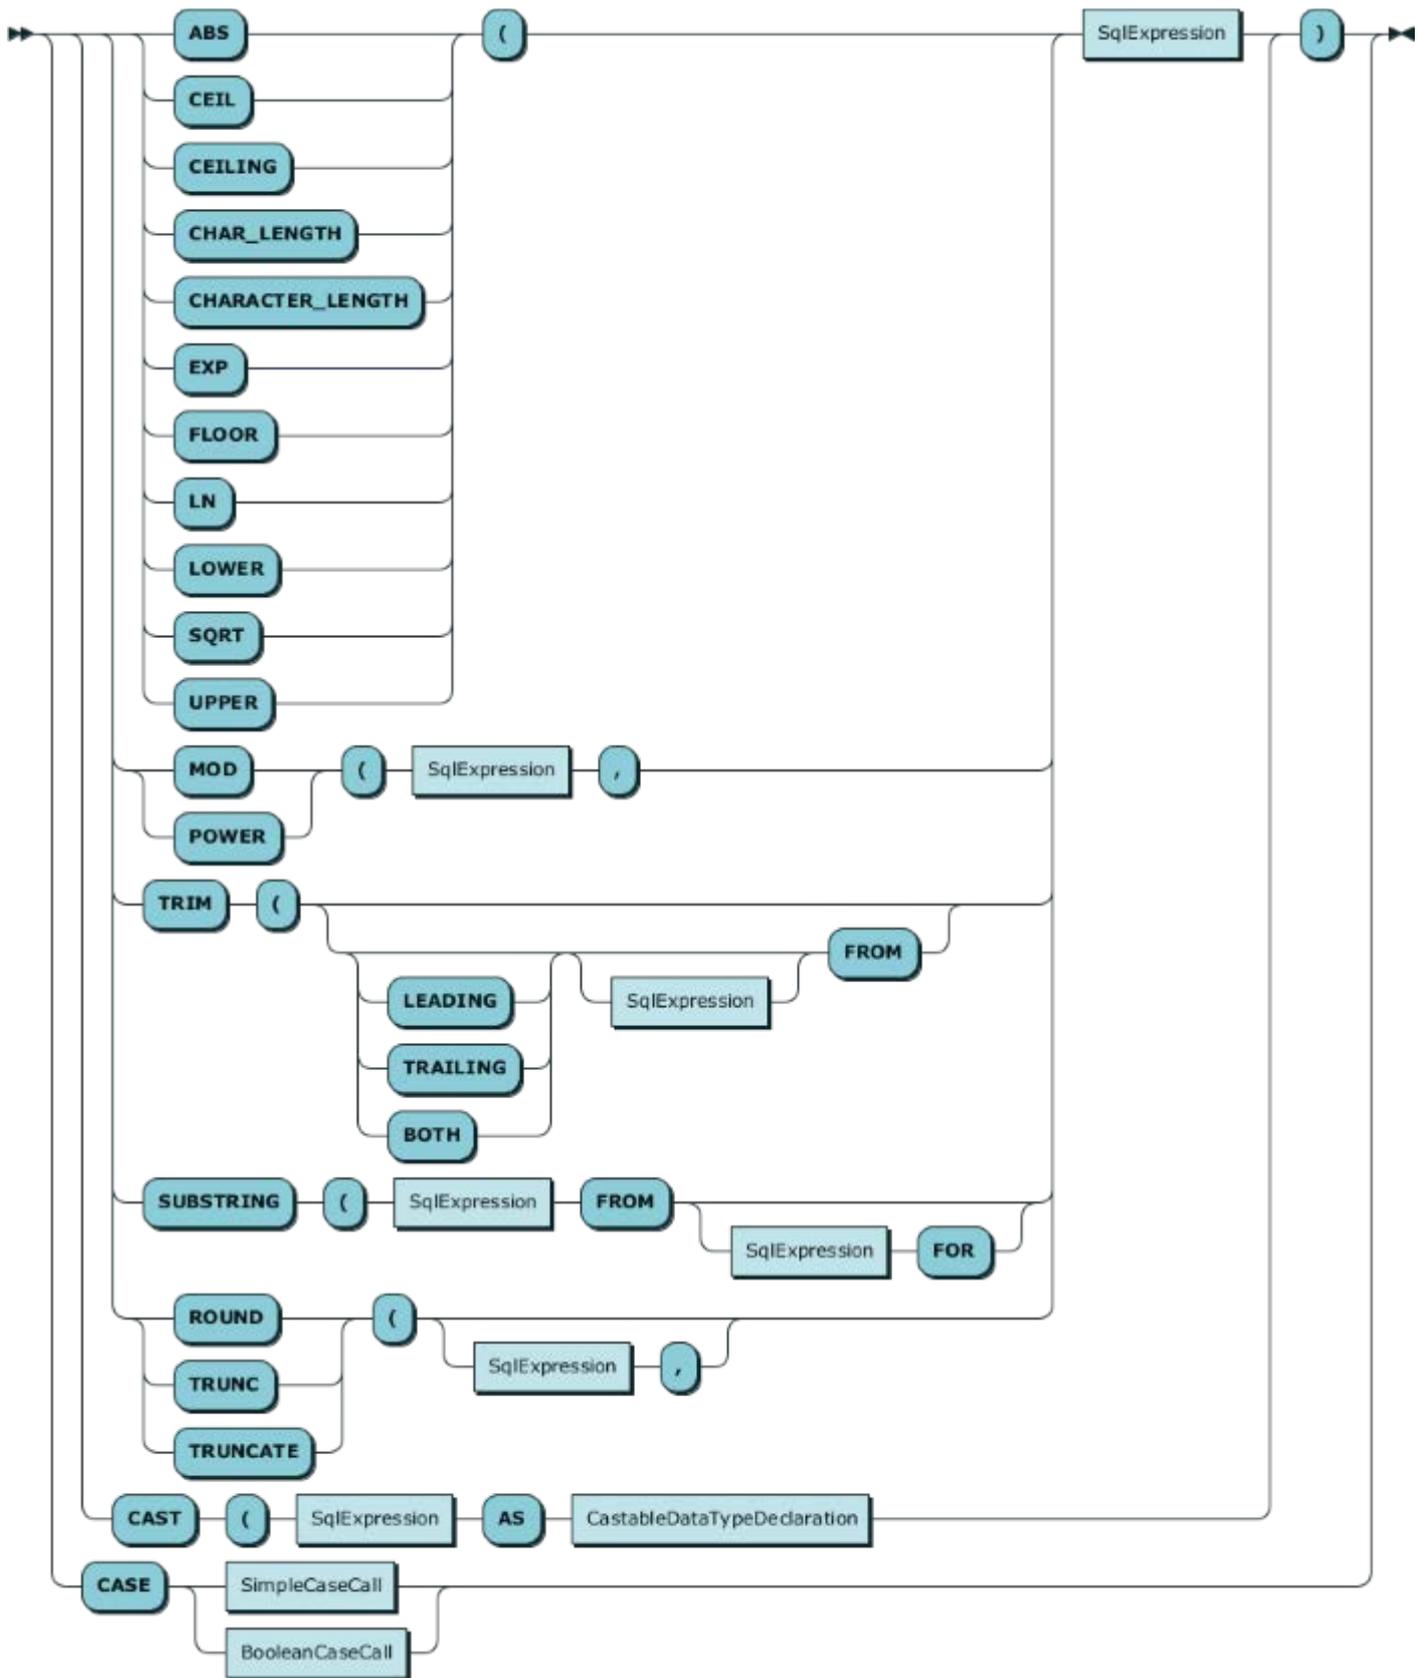
Definition:

Constant ::= 'NULL' | ('+' | '-')? S_NUMBER | S_CHAR_LITERAL | 'TRUE' | 'FALSE' | S_PARAMETER

Used In:

- ConstantList (page 386)
- SqlPrimaryExpression (page 373)
- WindowFunctionCall (page 389)

9.3.3.2.30 FunctionCall



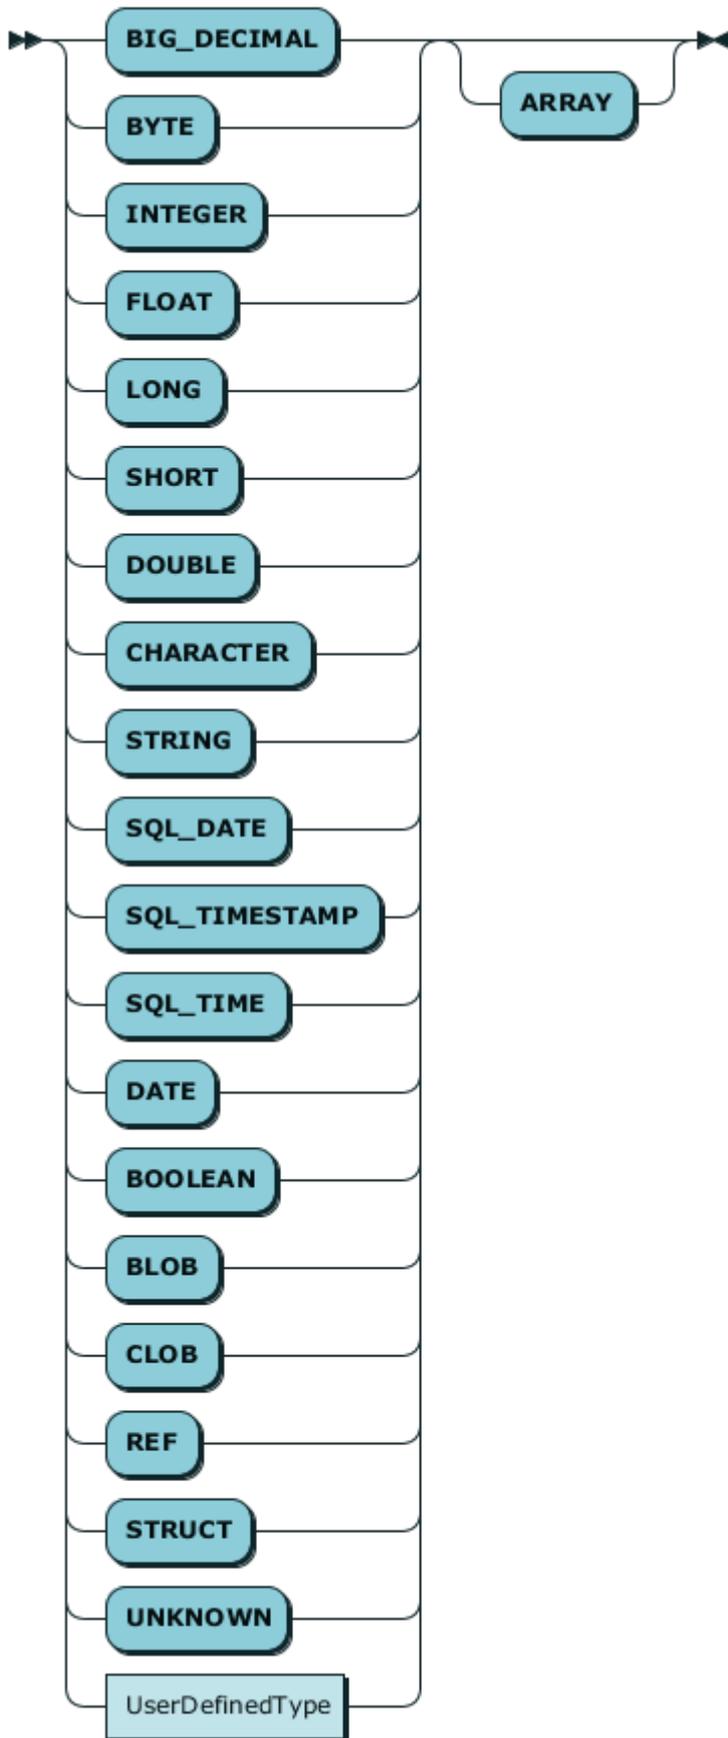
Definition:

FunctionCall ::= ((('ABS' | 'CEIL' | 'CEILING' | 'CHAR_LENGTH' | 'CHARACTER_LENGTH' | 'EXP' | 'FLOOR' | 'LN' | 'LOWER' | 'SQRT' | 'UPPER') '(' | ('MOD' | 'POWER') '(' SqlExpression ',' | 'TRIM' '(' (('LEADING' | 'TRAILING' | 'BOTH')? SqlExpression? 'FROM')? | 'SUBSTRING' '(' SqlExpression 'FROM' (SqlExpression 'FOR')? | ('ROUND' | 'TRUNC' | 'TRUNCATE') '(' (SqlExpression ',')?) SqlExpression | 'CAST' '(' SqlExpression 'AS' CastableDataTypeDeclaration) ')' | 'CASE' (SimpleCaseCall | BooleanCaseCall)

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.31 CastableDataTypeDeclaration



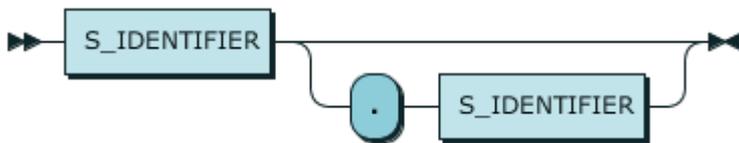
Definition:

CastableDataTypeDeclaration ::= ('BIG_DECIMAL' | 'BYTE' | 'INTEGER' | 'FLOAT' | 'LONG' | 'SHORT' | 'DOUBLE' | 'CHARACTER' | 'STRING' | 'SQL_DATE' | 'SQL_TIMESTAMP' | 'SQL_TIME' | 'DATE' | 'BOOLEAN' | 'BLOB' | 'CLOB' | 'REF' | 'STRUCT' | 'UNKNOWN' | UserDefinedType) 'ARRAY'?

Used In:

- FunctionCall (page 379)

9.3.3.2.32 UserDefinedType



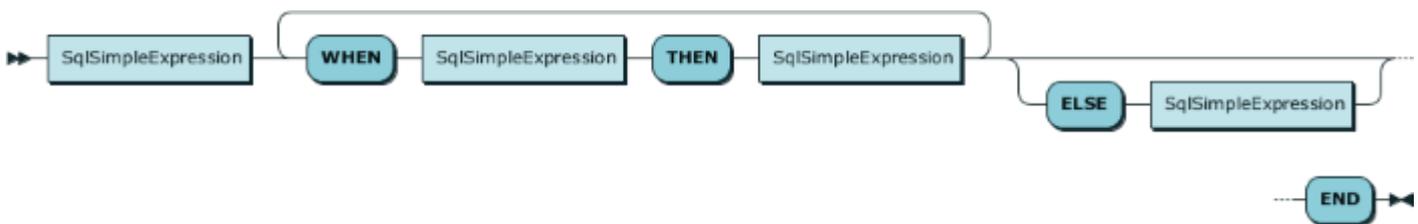
Definition:

UserDefinedType ::= S_IDENTIFIER ('.' S_IDENTIFIER)?

Used In:

- CastableDataTypeDeclaration (page 381)

9.3.3.2.33 SimpleCaseCall



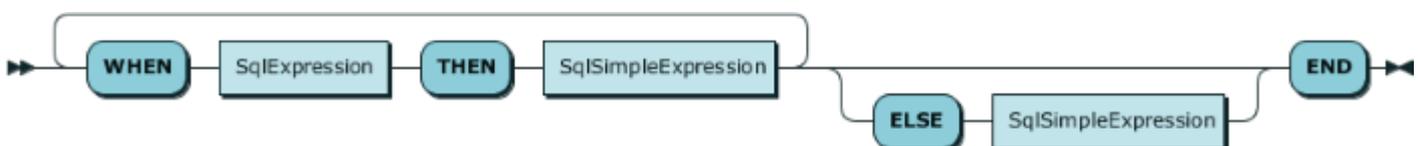
Definition:

SimpleCaseCall ::= SqlSimpleExpression ('WHEN' SqlSimpleExpression 'THEN' SqlSimpleExpression)+ ('ELSE' SqlSimpleExpression)? 'END'

Used In:

- FunctionCall (page 379)

9.3.3.2.34 BooleanCaseCall



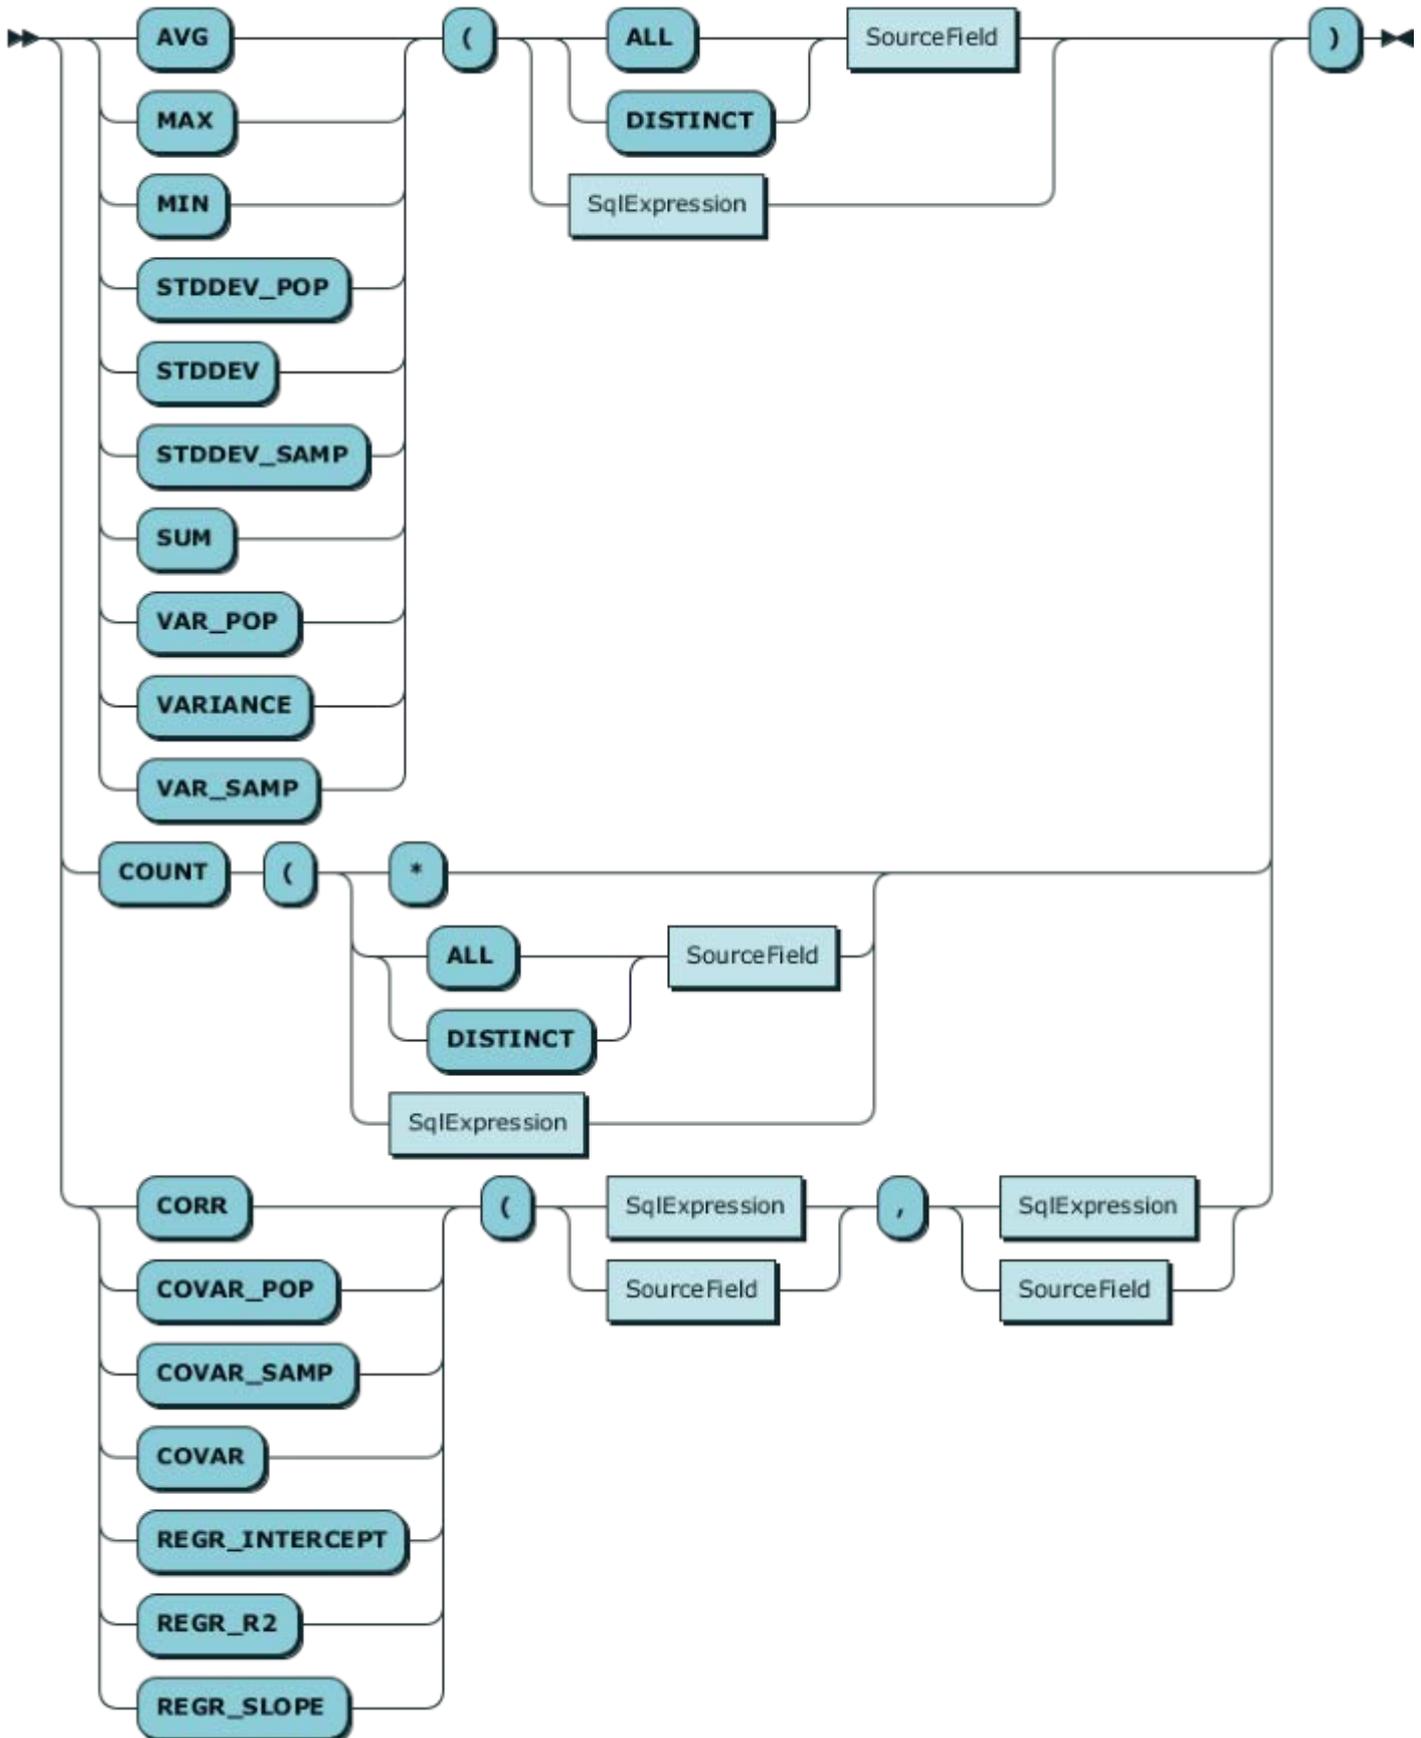
Definition:

BooleanCaseCall ::= ('WHEN' SqlExpression 'THEN' SqlSimpleExpression)+ ('ELSE' SqlSimpleExpression)? 'END'

Used In:

- [FunctionCall](#) (page 379)

9.3.3.2.35 AggregateFunctionCall



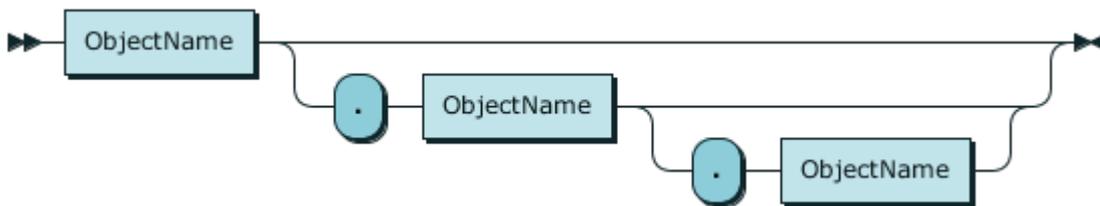
Definition:

AggregateFunctionCall ::= (('AVG' | 'MAX' | 'MIN' | 'STDDEV_POP' | 'STDDEV' | 'STDDEV_SAMP' | 'SUM' | 'VAR_POP' | 'VARIANCE' | 'VAR_SAMP') '(' (('ALL' | 'DISTINCT') SourceField | SqlExpression) | 'COUNT' '(' ('*' | ('ALL' | 'DISTINCT') SourceField | SqlExpression) | ('CORR' | 'COVAR_POP' | 'COVAR_SAMP' | 'COVAR' | 'REGR_INTERCEPT' | 'REGR_R2' | 'REGR_SLOPE') '(' (SqlExpression | SourceField) ',' (SqlExpression | SourceField)))'

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.36 SourceField



Definition:

SourceField ::= ObjectName ('.' ObjectName ('.' ObjectName)?)?

Used In:

- AggregateFunctionCall (page 384)
- SourceFieldList (page 393)
- WindowFunctionCall (page 389)

9.3.3.2.37 GenericFunctionCall



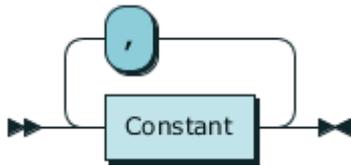
Definition:

GenericFunctionCall ::= ObjectName '(' SqlExpressionList? (';' ConstantList)? ')'

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.38 ConstantList



Definition:

ConstantList ::= Constant (',' Constant)*

Used In:

- GenericFunctionCall (page 385)

9.3.3.2.39 WindowSpecification



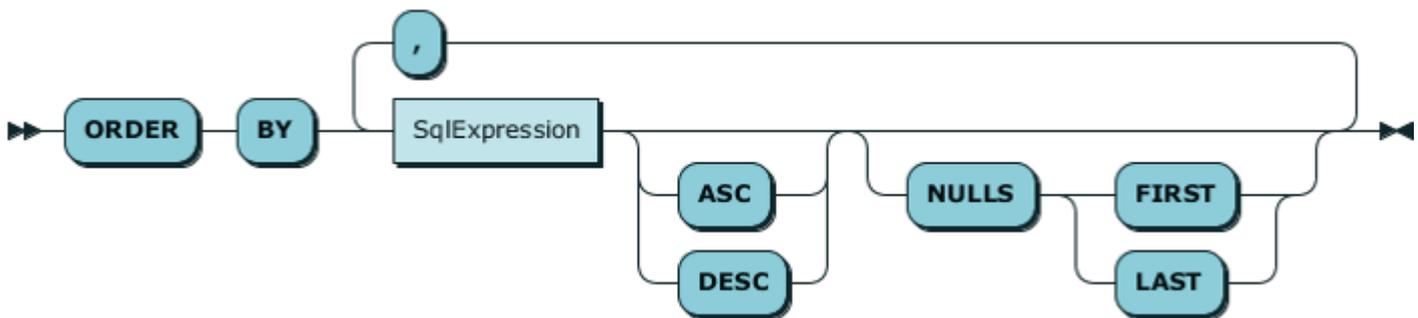
Definition:

WindowSpecification ::= 'OVER' '(' ('PARTITION' 'BY' SqlExpressionList)? OrderByClause? WindowFrameClause? ')'

Used In:

- SqlPrimaryExpression (page 373)
- WindowFunctionCall (page 389)

9.3.3.2.40 OrderByClause



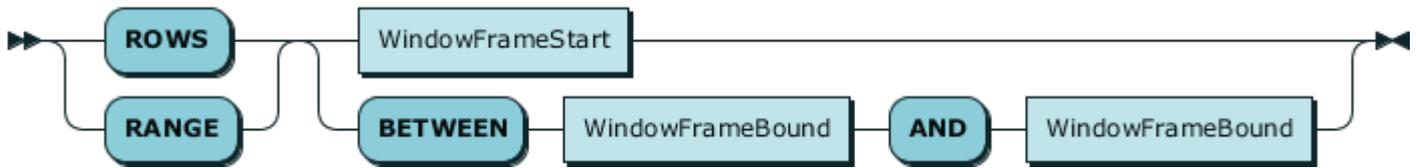
Definition:

OrderByClause ::= 'ORDER' 'BY' SqlExpression ('ASC' | 'DESC')? ('NULLS' ('FIRST' | 'LAST'))? (',' SqlExpression ('ASC' | 'DESC')? ('NULLS' ('FIRST' | 'LAST'))?)?*

Used In:

- SelectWithOrder (page 368)
- WindowSpecification (page 386)

9.3.3.2.41 WindowFrameClause



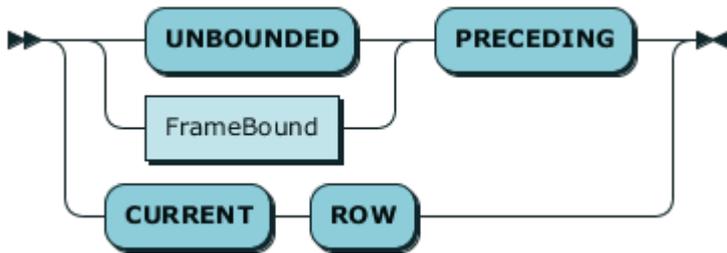
Definition:

WindowFrameClause ::= ('ROWS' | 'RANGE') (WindowFrameStart | 'BETWEEN' WindowFrameBound 'AND' WindowFrameBound)

Used In:

- WindowSpecification (page 386)

9.3.3.2.42 WindowFrameStart



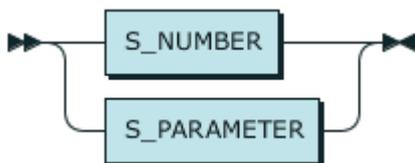
Definition:

WindowFrameStart ::= ('UNBOUNDED' | FrameBound) 'PRECEDING' | 'CURRENT' 'ROW'

Used In:

- WindowFrameBound (page 388)
- WindowFrameClause (page 387)

9.3.3.2.43 FrameBound



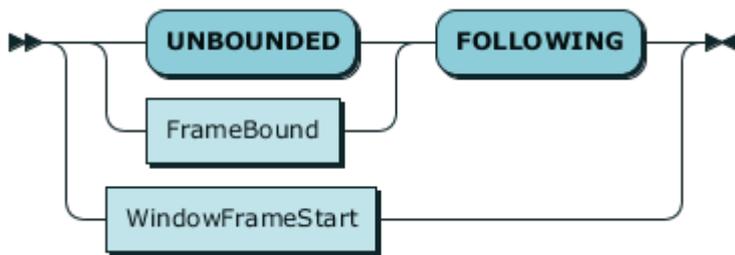
Definition:

FrameBound ::= S_NUMBER | S_PARAMETER

Used In:

- WindowFrameBound (page 388)
- WindowFrameStart (page 387)

9.3.3.2.44 WindowFrameBound



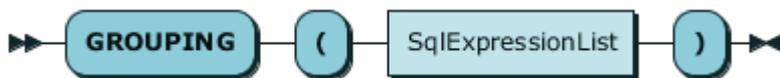
Definition:

WindowFrameBound ::= ('UNBOUNDED' | FrameBound) 'FOLLOWING' | WindowFrameStart

Used In:

- WindowFrameClause (page 387)

9.3.3.2.45 GroupingOperation



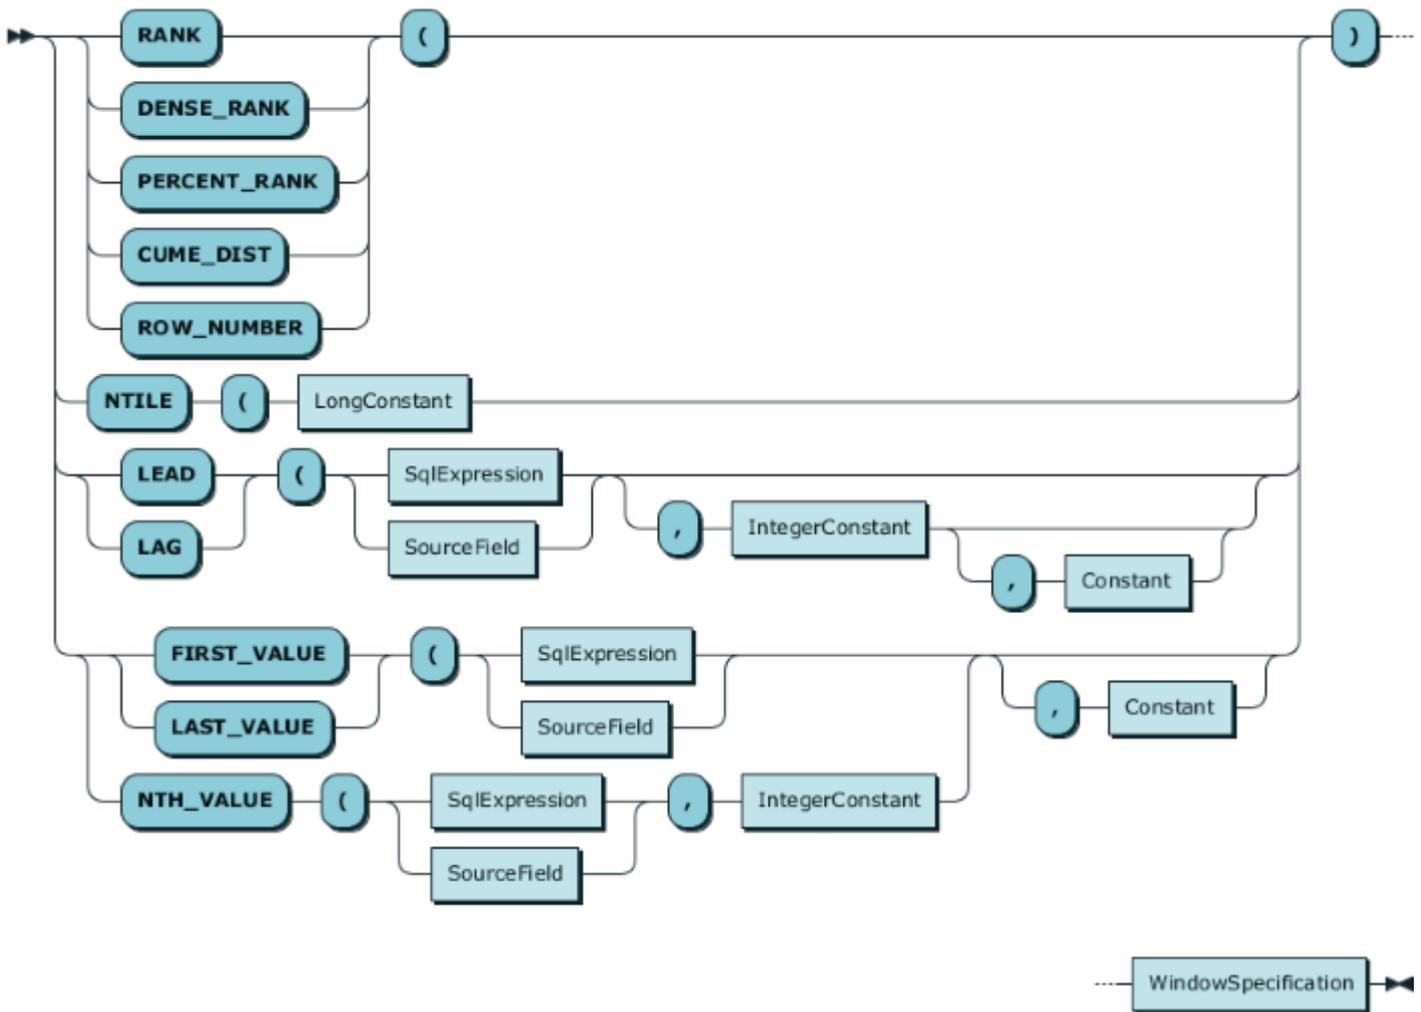
Definition:

GroupingOperation ::= 'GROUPING' '(' SqlExpressionList ')'

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.46 WindowFunctionCall



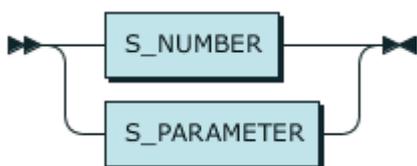
Definition:

WindowFunctionCall ::= (('RANK' | 'DENSE_RANK' | 'PERCENT_RANK' | 'CUME_DIST' | 'ROW_NUMBER') '(' | 'NTILE' '(' LongConstant | ('LEAD' | 'LAG') '(' (SqlExpression | SourceField) (',' IntegerConstant (',' Constant)?)? | (('FIRST_VALUE' | 'LAST_VALUE') '(' (SqlExpression | SourceField) | 'NTH_VALUE' '(' (SqlExpression | SourceField) ',' IntegerConstant (',' Constant)?))' WindowSpecification

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.47 LongConstant



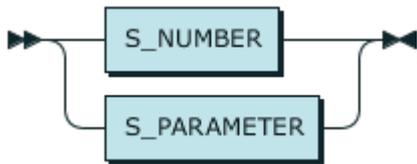
Definition:

LongConstant ::= S_NUMBER | S_PARAMETER

Used In:

- LimitClause (page 398)
- WindowFunctionCall (page 389)

9.3.3.2.48 IntegerConstant



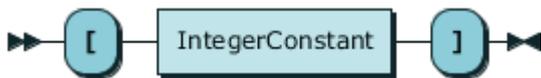
Definition:

IntegerConstant ::= S_NUMBER | S_PARAMETER

Used In:

- ArrayElementReference (page 390)
- WindowFunctionCall (page 389)

9.3.3.2.49 ArrayElementReference



Definition:

ArrayElementReference ::= '[' IntegerConstant ']'

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.50 ArrayValueConstructor



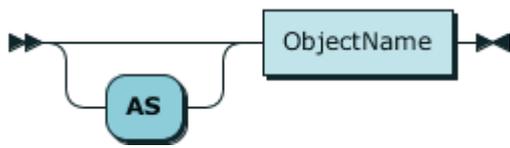
Definition:

ArrayValueConstructor ::= 'ARRAY' '[' SqlExpressionList ']'

Used In:

- SqlPrimaryExpression (page 373)

9.3.3.2.51 Alias



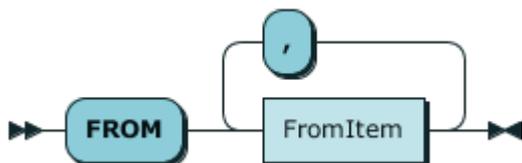
Definition:

Alias ::= 'AS'? ObjectName

Used In:

- SelectItem (page 370)
- SourceReferenceOrSubquery (page 391)

9.3.3.2.52 FromClause



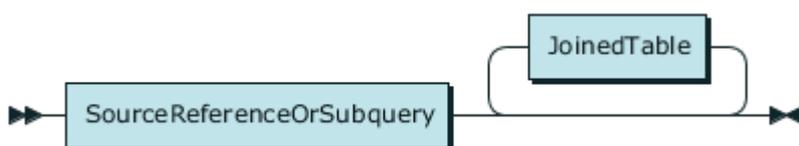
Definition:

FromClause ::= 'FROM' FromItem (',' FromItem)*

Used In:

- SelectWithoutOrder (page 369)

9.3.3.2.53 FromItem



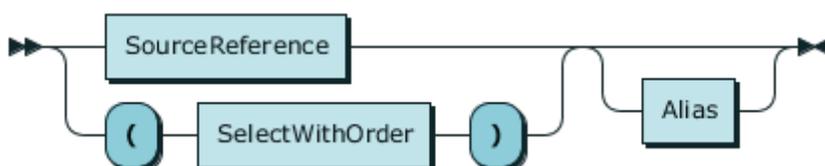
Definition:

FromItem ::= SourceReferenceOrSubquery JoinedTable*

Used In:

- FromClause (page 391)

9.3.3.2.54 SourceReferenceOrSubquery



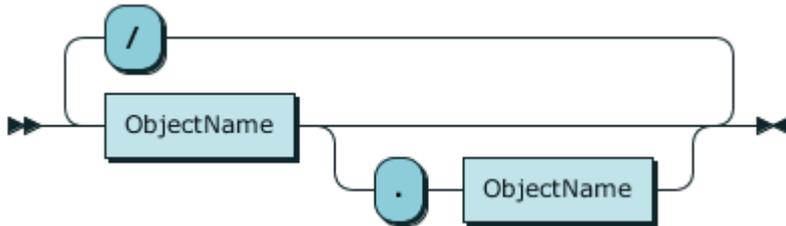
Definition:

SourceReferenceOrSubquery ::= (SourceReference | '(' SelectWithOrder ')') Alias?

Used In:

- FromItem (page 391)
- JoinedTable (page 392)

9.3.3.2.55 SourceReference



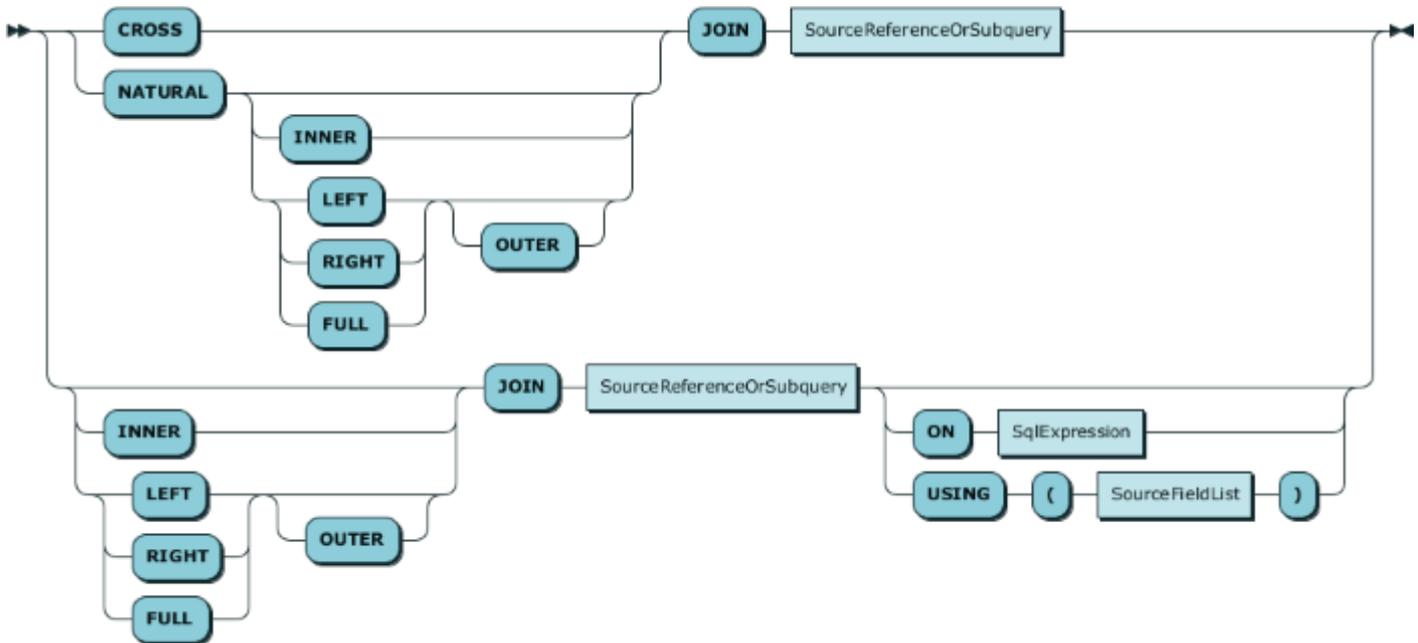
Definition:

SourceReference ::= ObjectName ('.' ObjectName)? ('/' ObjectName ('.' ObjectName)?)*

Used In:

- SourceReferenceOrSubquery (page 391)

9.3.3.2.56 JoinedTable



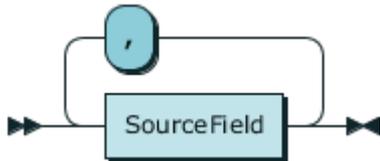
Definition:

JoinedTable ::= ('CROSS' | 'NATURAL' ('INNER' | ('LEFT' | 'RIGHT' | 'FULL') 'OUTER'?)) 'JOIN'
 SourceReferenceOrSubquery | ('INNER' | ('LEFT' | 'RIGHT' | 'FULL') 'OUTER'?) 'JOIN'
 SourceReferenceOrSubquery ('ON' SqlExpression | 'USING' '(' SourceFieldList ')') ?

Used In:

- FromItem (page 391)

9.3.3.2.57 SourceFieldList



Definition:

SourceFieldList ::= SourceField (',' SourceField) *

Used In:

- JoinedTable (page 392)

9.3.3.2.58 WhereClause



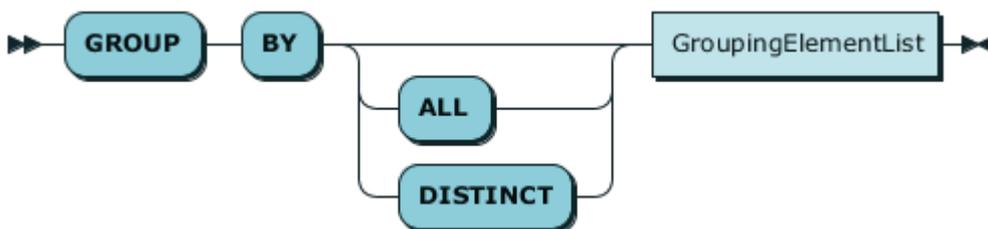
Definition:

WhereClause ::= 'WHERE' SqlExpression

Used In:

- SelectWithoutOrder (page 369)

9.3.3.2.59 GroupByClause



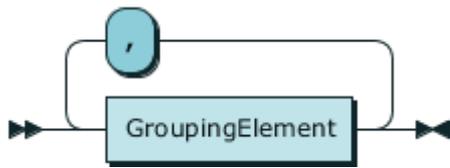
Definition:

GroupByClause ::= 'GROUP' 'BY' ('ALL' | 'DISTINCT') ? GroupingElementList

Used In:

- SelectWithoutOrder (page 369)

9.3.3.2.60 GroupingElementList



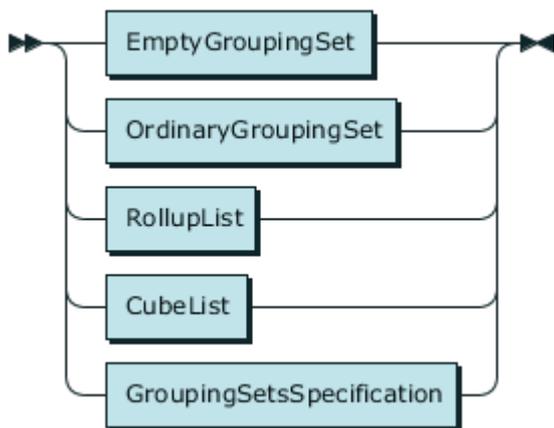
Definition:

GroupingElementList ::= GroupingElement (',' GroupingElement)*

Used In:

- GroupByClause (page 393)

9.3.3.2.61 GroupingElement



Definition:

GroupingElement ::= EmptyGroupingSet | OrdinaryGroupingSet | RollupList | CubeList | GroupingSetsSpecification

Used In:

- GroupingElementList (page 394)

9.3.3.2.62 EmptyGroupingSet



Definition:

EmptyGroupingSet ::= '(' ')'

Used In:

- GroupingElement (page 394)

- GroupingSet (page 397)

9.3.3.2.63 OrdinaryGroupingSet



Definition:

OrdinaryGroupingSet ::= '(' GroupingExpressionReferenceList ')' | GroupingExpressionReference

Used In:

- GroupingElement (page 394)
- GroupingSet (page 397)
- OrdinaryGroupingSetList (page 396)

9.3.3.2.64 GroupingExpressionReferenceList



Definition:

GroupingExpressionReferenceList ::= GroupingExpressionReference (',' GroupingExpressionReference)*

Used In:

- OrdinaryGroupingSet (page 395)

9.3.3.2.65 GroupingExpressionReference



Definition:

GroupingExpressionReference ::= SqlExpression

Used In:

- GroupingExpressionReferenceList (page 395)
- OrdinaryGroupingSet (page 395)

9.3.3.2.66 RollupList



Definition:

RollupList ::= 'ROLLUP' '(' OrdinaryGroupingSetList ')'

Used In:

- GroupingElement (page 394)
- GroupingSet (page 397)

9.3.3.2.67 OrdinaryGroupingSetList



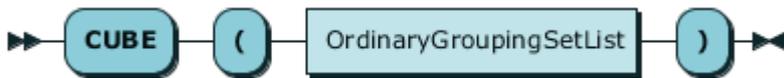
Definition:

OrdinaryGroupingSetList ::= OrdinaryGroupingSet (',' OrdinaryGroupingSet)*

Used In:

- CubeList (page 396)
- RollupList (page 396)

9.3.3.2.68 CubeList



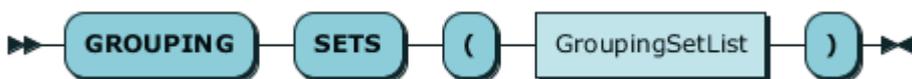
Definition:

CubeList ::= 'CUBE' '(' OrdinaryGroupingSetList ')'

Used In:

- GroupingElement (page 394)
- GroupingSet (page 397)

9.3.3.2.69 GroupingSetsSpecification



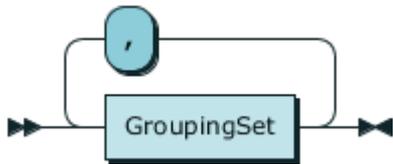
Definition:

GroupingSetsSpecification ::= 'GROUPING' 'SETS' '(' GroupingSetList ')'

Used In:

- GroupingElement (page 394)
- GroupingSet (page 397)

9.3.3.2.70 GroupingSetList



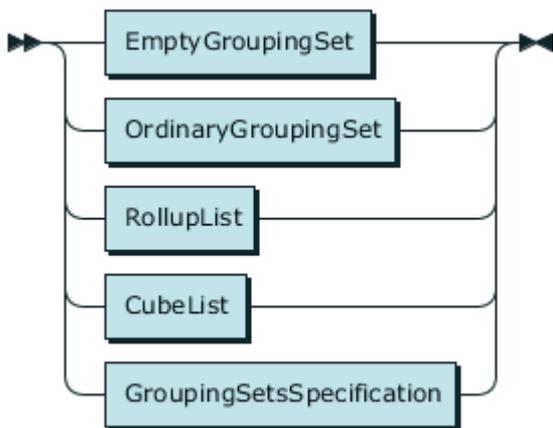
Definition:

GroupingSetList ::= GroupingSet (',' GroupingSet)*

Used In:

- GroupingSetsSpecification (page 396)

9.3.3.2.71 GroupingSet



Definition:

GroupingSet ::= EmptyGroupingSet | OrdinaryGroupingSet | RollupList | CubeList | GroupingSetsSpecification

Used In:

- GroupingSetList (page 397)

9.3.3.2.72 HavingClause



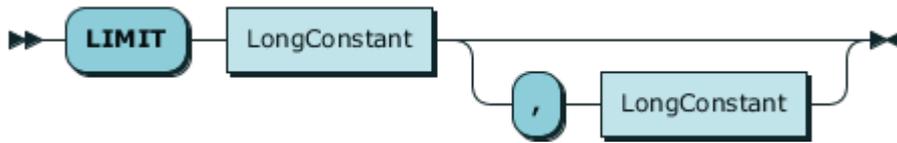
Definition:

HavingClause ::= 'HAVING' SqlExpression

Used In:

- SelectWithoutOrder (page 369)

9.3.3.2.73 LimitClause



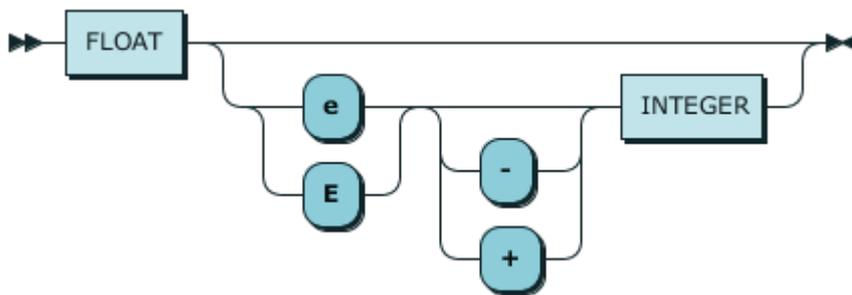
Definition:

LimitClause ::= 'LIMIT' LongConstant (',' LongConstant)?

Used In:

- SelectWithOrder (page 369)

9.3.3.2.74 S_NUMBER



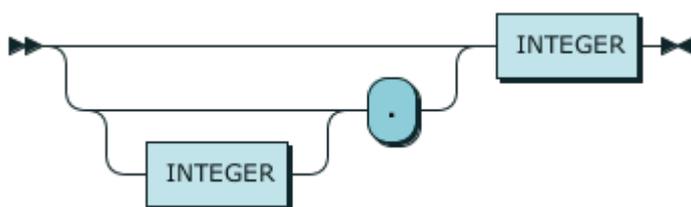
Definition:

S_NUMBER ::= FLOAT ([eE] [-+]? INTEGER)?

Used In:

- Constant (page 378)
- FrameBound (page 387)
- IntegerConstant (page 390)
- LongConstant (page 389)

9.3.3.2.75 FLOAT



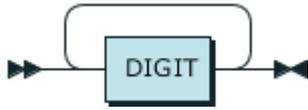
Definition:

FLOAT ::= (INTEGER? '.')? INTEGER

Used In:

- S_NUMBER (page 398)

9.3.3.2.76 INTEGER



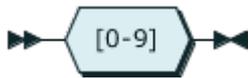
Definition:

INTEGER ::= DIGIT+

Used In:

- FLOAT (page 398)
- S_NUMBER (page 398)

9.3.3.2.77 DIGIT



Definition:

DIGIT ::= [0-9]

Used In:

- INTEGER (page 399)

9.3.3.2.78 S_PARAMETER



Definition:

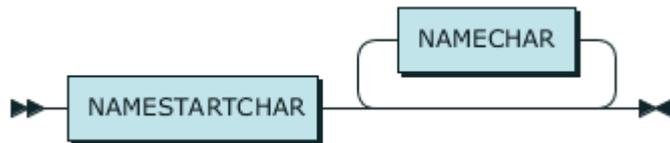
S_PARAMETER ::= ':' (S_IDENTIFIER | S_QUOTED_IDENTIFIER)

Used In:

- CharacterConstant (page 367)
- Constant (page 378)
- FrameBound (page 387)

- IntegerConstant (page 390)
- LongConstant (page 389)

9.3.3.2.79 S_IDENTIFIER



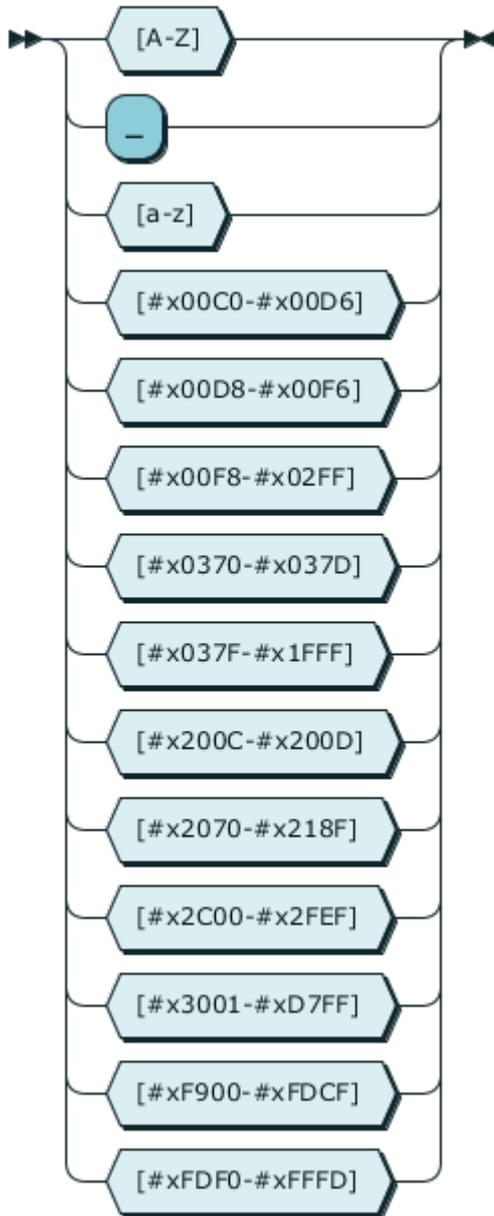
Definition:

S_IDENTIFIER ::= NAMESTARTCHAR NAMECHAR*

Used In:

- ObjectName (page 371)
- S_PARAMETER (page 399)
- UserDefinedType (page 382)

9.3.3.2.80 NAMESTARTCHAR



Definition:

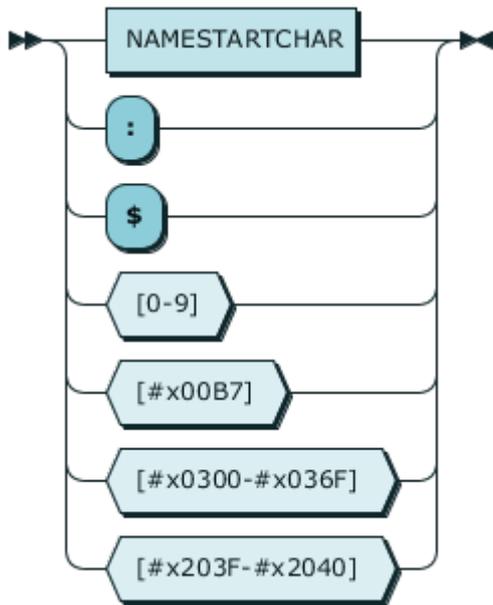
NAMESTARTCHAR ::=

[A-Z_a-z#x00C0-#x00D6#x00D8-#x00F6#x00F8-#x02FF#x0370-#x037D#x037F-#x1FFF#x200C-#x200D#x2070-#x218F#x2C00-#x2FEF#x3001-#xD7FF#xF900-#xFDCF#xFDF0-#xFFFD]

Used In:

- NAMECHAR (page 402)
- S_IDENTIFIER (page 400)

9.3.3.2.81 NAMECHAR



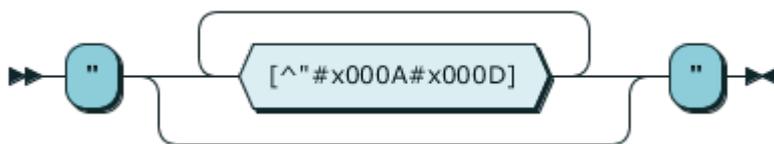
Definition:

NAMECHAR ::= NAMESTARTCHAR | [:\$0-9#x00B7#x0300-#x036F#x203F-#x2040]

Used In:

- S_IDENTIFIER (page 400)

9.3.3.2.82 S_QUOTED_IDENTIFIER



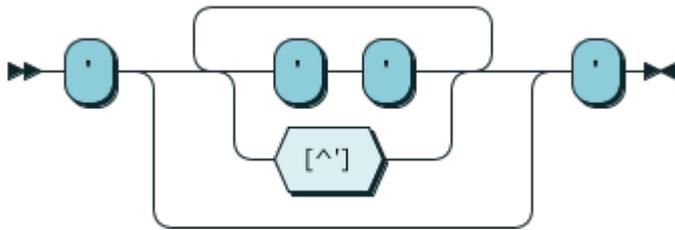
Definition:

S_QUOTED_IDENTIFIER ::= '"' [^" #x000A #x000D]* '''

Used In:

- ObjectName (page 371)
- S_PARAMETER (page 399)

9.3.3.2.83 S_CHAR_LITERAL



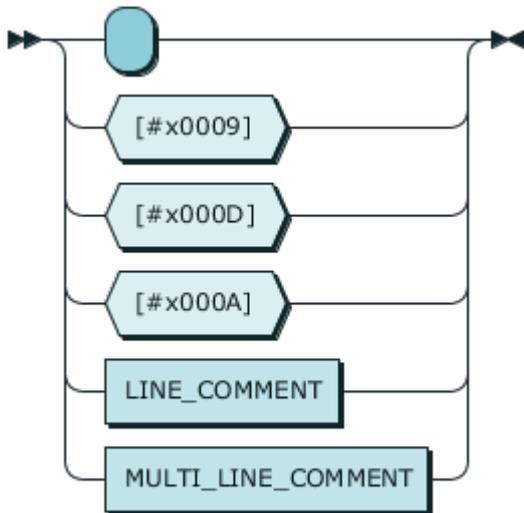
Definition:

S_CHAR_LITERAL ::= `'' ('' '' | [^']) * ''`

Used In:

- CharacterConstant (page 367)
- Constant (page 378)

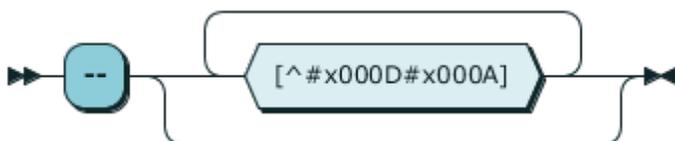
9.3.3.2.84 IgnorableWhitespace



Definition:

IgnorableWhitespace ::= `[#x0009#x000D#x000A] | LINE_COMMENT | MULTI_LINE_COMMENT`

9.3.3.2.85 LINE_COMMENT



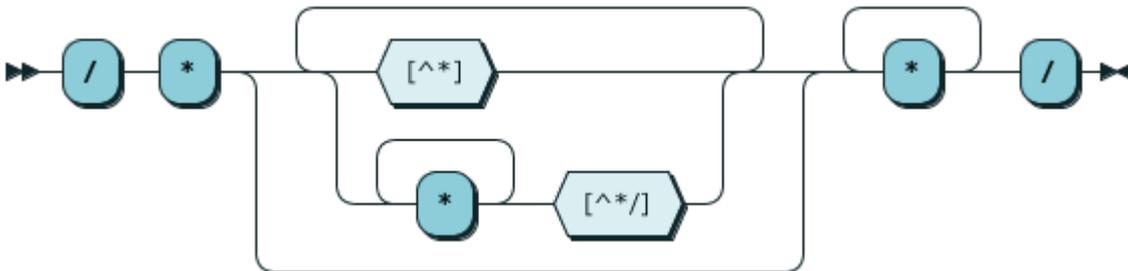
Definition:

LINE_COMMENT ::= '--' [^#x000D#x000A]*

Used In:

- IgnorableWhitespace (page 403)

9.3.3.2.86 MULTI_LINE_COMMENT



Definition:

MULTI_LINE_COMMENT ::= '/' '*' ([^*] | '*' + [^*/])* '*' + '/'

Used In:

- IgnorableWhitespace (page 403)

9.3.4 RAQL extensions

MashZone NextGen provides both plain functions as well as analytic functions that you may use in RAQL queries. While some extensions are based on standard SQL definitions, others are RAQL-specific extensions that widens the support of analytic tasks you can tackle. You can also define and add your own plain or analytic functions to RAQL. See [Creating new user-defined functions](#) (page 419) for instructions.

Note that in contrast to the case-insensitive handling of SQL functions, RAQL-specific functions are handled case-sensitively.

9.3.4.1 RAQL function extensions

The following function extensions are RAQL-specific and are not part of the SQL standard. For a discussion of the functions defined in SQL, we refer to the SQL documentation. See also section [FunctionCall](#) (page 379) for more details on the standard SQL functions being supported.

Note: Some function extensions from version 3.7 are deprecated and no longer supported in versions 3.8 and onward. See [8.3.4.1.4 Migrate RAQL queries from Version 3.7 to 3.8](#) (page 416) for a complete list and the alternate functions from version 3.8 to use as replacements.

9.3.4.1.1 Plain functions

Plain functions can be used in SELECT, OVER, WHERE, ORDER BY, GROUP BY and HAVING clauses in RAQL queries. They typically either change the data type of column values, extract part of the values, or transform values in some way.

Plain functions are applied individually to values of one row without access to values in other rows and they return a single value.

Function	Description
column_name(Any col_1, Any col_2, ..., Any col_n, Number index) column_name(Any col)	Returns the column value at position index from the columns col_1 , col_2 , ..., col_n . The index will be rounded if not an integer value and is not required if only one column is used.
column_type(Any col_1, Any col_2, ..., Any col_n, Number index) column_type(Any col)	Returns the column name at position index from the columns col_1 , col_2 , ..., col_n . The index will be rounded if not an integer value and is not required if only one column is used.
custom_month(String col-or-expr)	Extracts the month from dates in the format EEE MMM dd HH:mm:ss z yyyy and returns the number of the month as an integer. This is based on English month names.
dayOfWeek(Date col-or-expr) dayOfWeek(String col-or-expr) dayOfWeek(Long col-or-expr)	Converts the values for this column or expression to a date and returns the day of the week as an integer [1 - 7]. This number is specific to the locale for the MashZone NextGen Server. For example, day 1 is Sunday in the United States but Monday in France.
epoc_sec_to_date	Represents internally the Java notation of time as milliseconds since midnight, January 1, 1970 UTC. <ul style="list-style-type: none"> ▪ Conversion of seconds (not milliseconds!) since above reference time to date ▪ Three variants of function using either string, double, or long as input Examples <ul style="list-style-type: none"> ▪ epoc_sec_to_date('123456') ▪ epoc_sec_to_date(123456.00) ▪ epoc_sec_to_date(123456)
extract_date(Date dt, String field)	Extracts the portion of the date or time from values for this column or expression that is identified by the dt parameter. Valid field values are: <ul style="list-style-type: none"> ▪ day

Function	Description
	<ul style="list-style-type: none"> ▪ day-of-week ▪ hour ▪ minute ▪ month ▪ monthname ▪ quarter ▪ second ▪ week ▪ year
<p>extract_hour(Date dt) extract_hour(String dt) extract_hour(Long dt)</p>	<p>Extracts the hour portion of the date and time in the column or expression that is identified by the dt parameter.</p>
<p>extract_minute(Date dt) extract_minute(String dt) extract_minute(Long dt)</p>	<p>Extracts the minute portion of the date and time in the column or expression that is identified by the dt parameter.</p>
<p>extract_month(Date dt) extract_month(String dt) extract_month(Long dt)</p>	<p>Extracts the month portion of the date and time in the column or expression that is identified by the dt parameter.</p>
<p>extract_second(Date dt) extract_second(String dt) extract_second(Long dt)</p>	<p>Extracts the second portion of the date and time in the column or expression that is identified by the dt parameter.</p>
<p>extract_week(Date dt) extract_week(String dt) extract_week(Long dt)</p>	<p>Extracts the week portion of the date and time in the column or expression that is identified by the dt parameter.</p>
<p>extract_year(Date dt) extract_year(String dt) extract_year(Long dt)</p>	<p>Extracts the year portion of the date and time in the column or expression that is identified by the dt parameter.</p>
<p>format_date(dateColumn, datePattern, TimeZone)</p>	<p>Represents string of date values in target format and target timezone.</p> <ul style="list-style-type: none"> ▪ datePattern may be 'short', 'medium', 'long', 'full' or any date pattern as specified in Java's SimpleDateFormat. ▪ TimeZone may be any time zone as specified in Java's SimpleDateFormat <p>Examples</p>

Function	Description
	<ul style="list-style-type: none"> ▪ <code>format_date (dateColumn , ' yyyy -MM- dd ', 'GMT')</code> ▪ <code>format_date (dateColumn , ' h:mm a ', 'PST-08:00')</code>
<code>format_date (dateColumn, datePattern)</code>	Corresponds to <code>format_date(dateColumn, datePattern, 'GMT')</code>
<code>format_date (dateColumn)</code>	Corresponds to <code>format_date(dateColumn, 'yyyy-MM-dd'T'HH:mm:ss.SSSZ', 'GMT')</code>
<code>geo_distance(String column-or-lat1, String column-or-long1, String column-or-lat2, String column-or-long2)</code>	Calculates the geographical distance between two locations identified by latitude and longitude. Commonly, one set of coordinates is provided from two columns or expressions in the data and the other is provided as literal values. If any parameter value is null, this returns a null value.
<code>matches(String col-or-expr, String pattern)</code>	<p>Determines if a string matching the regular expression in pattern exists anywhere within the values of the column or expression. Matching is case insensitive. Returns a boolean value (true or false).</p> <p>For example:</p> <p>where matches(firstname, 'ra[n l]')='true'</p> <p>Will match Frank, Gerald or Randy in the firstname column. To perform exact regular expression matches, with no implicit wildcards, see the regex function.</p>
<code>nvl(String col-or-expr, String str)</code> <code>nvl(Date col-or-expr, String str)</code> <code>nvl(Number col-or-expr, String str)</code>	<p>Replaces null values in the specified column or expression with the specified string value.</p> <p>Note: In previous releases, string columns with only white space characters (space, tab, etc.) were treated as a null value. Effective in 3.8, columns with string values are considered to be null only if they have no value or they contain an empty string.</p>
<code>project(Xcol_1, X col_2, ..., X col_n, Number index)</code>	Returns the value at position index from the columns col_1, col_2, ..., col_n of a common type X . The index will be rounded if not an integer value.
<code>quarter(Date col-or-expr)</code> <code>quarter(String col-or-expr)</code> <code>quarter(Long col-or-expr)</code>	Converts the values for this column or expression to a date and returns the number of the quarter as an integer.
<code>regex(String col-or-expr, String</code>	Determines if a string matching the regular expression in pattern exists in the values of the column or expression as

Function	Description
<p>pattern)</p>	<p>specified. Matching is case insensitive. Returns a boolean value (true or false).</p> <p>Note: This function does not add any implicit wildcards to the regular expression. To find a string anywhere within column values, use the matches function.</p> <p>For example:</p> <p>where regex(firstname,'ra[n l]')=true</p> <p>Will match Ralph or Randy in the firstname column, but will not match Frank or Gerald.</p>
<p>split_part(String col-or-expr, String delim, Integer part)</p> <p>split_part(String col-or-expr, Character delim, Integer part)</p>	<p>Returns the nth part of the values of this column or expression when each value is split into parts at each delimiter defined in delim. For example:</p> <p>split_part("INV:2012:GHI 345",":",3)</p> <p>Returns the string "GHI345".</p>
<p>split_regex((String col-or-expr, String pattern, Integer part)</p>	<p>This function is identical to split_part with the exception that the delimiter used to split the column values from pattern is a regular expression.</p>
<p>time_mask(Date column, String mask)</p> <p>time_mask(String column, String mask)</p>	<p>Converts the value for this column or expression to a date and time, if needed, in the form yyyy-MM-dd:HH:mm:ss.</p> <p>Note: If the values for the column do not contain times, the time is set to 00:00:00.</p> <p>It extracts the portion of this date and time through the time part identified by the mask.</p> <p>Valid masks include:</p> <ul style="list-style-type: none"> ▪ y = returns only the year. ▪ M = returns the year and month. ▪ d = returns the year, month and day. ▪ H = returns the full date and hour. ▪ m = returns the full date, hour and minute. ▪ s = returns the full date and time.
<p>to_date(String col-or-expr)</p> <p>to_date(String col-or-expr, String format)</p>	<p>Casts the values for this column or expression to Date.</p> <p>If no format pattern is provided, this function can have a negative impact on performance as it attempts to convert the string iterating over the following date formats:</p> <ul style="list-style-type: none"> ▪ M/dd/yy or the short format for the current locale ▪ MMM dd, yyyy or the long format for the current locale ▪ EEE MMM dd, yyyy or the full format for the current

Function	Description
	<p>locale</p> <ul style="list-style-type: none"> ▪ yyyy-MM-dd ▪ EEE MM dd HH:mm:ss z yyyy ▪ EEE, dd MMM yyyy HH:mm:ss Z ▪ yyyy-MM-DD'T'HH:mm:ss'Z' <p>Use the format parameter to define the date and time pattern used in values for this column.</p> <p>Note: For more information on date format patterns, see the Java SimpleDateFormat class.</p>
truncate(Number col-or-expr, Number decimals)	For numeric values or expressions, this truncates the number to the number of decimals places specified. If the decimals parameter is omitted, truncates the number to zero decimals.

9.3.4.1.2 Analytic functions

Analytic functions most commonly perform calculations using sets of rows within a dataset. This may be the entire dataset or specific sets of rows defined as groups, partitions or windows.

Aggregation analytic functions: use all rows in the dataset, group, partition or window to perform a calculation and return a single value. For example, sum adds the values of all rows in the current scope. Aggregate analytic functions can be used in the SELECT or HAVING clauses of RAQL queries:

- To return a single value for each group defined in a GROUP BY clause.
- To return a single value for each partition or each window defined in an OVER clause. This single value is added to each row in the partition or window. You can also use aggregate analytic functions to return running calculations for partitions or windows defined in an OVER clause.
- To return a single value for the entire dataset if no group definition or partition definition is specified.

Window analytic functions: use some of the rows in a partition or window to perform a calculation and return a value for each row. Typically, this uses rows that are relative to the current row, such as FIRST_VALUE or ROW_NUMBER.

You must include an OVER clause in queries that use window analytic functions. This defines the partitions or windows used by the function.

As with plain functions, RAQL provides analytic functions defined in the SQL standard as well as RAQL-specific extensions. See AggregateFunctionCall (page 384) for more details on the standard SQL aggregation functions being supported.

The RAQL-specific extensions are:

Aggregate	Window	Function	Description
✓		correlation(Number col-or-expr, Number col-or-expr)	Deprecated. See Update analytic functions for 3.8 (page 418) for alternatives.
✓		covar_samp(Number col-or-expr, Number col-or-expr)	Returns the sample covariance (an estimate) of the two specified columns. See covariance for more information.
	✓	analytics.discretize(Number col-or-expr, int binCount)	<p>This function segments a continuous range of values for a column or expression into discrete bins based on the number of bins specified. If the value:</p> <ul style="list-style-type: none"> ▪ \leq minimum value for the column, it returns 0. ▪ \geq maximum value for the column, it returns binCount - 1. ▪ $<$ maximum and $>$ minimum value for the column, it returns a fraction between 0 and binCount - 1 indicating the value's relative position within the number of bins. <p>So $0 \leq$ return value < 1, indicates the first bin, $1 \leq$ return value < 2, indicates the second bin, and so on.</p> <p>Note that this function does not support windows. It can only be applied to the entire partition.</p>
✓		gmean(Number col-or-expr)	Returns the geometric mean of the values for this column or expression within the partition or the current window in this

Aggregate	Window	Function	Description
			<p>partition.</p> <p>The geometric mean is typically used to define a mean when more than one property is involved, especially if the scale for the properties is different.</p>
✓		<p>analytics.kmeans_clusters(Number col-or-expr1[,Number col-or-expr2,...Number col-or-exprN]; Integer k, Integer iterations, String measure)</p>	<p>Returns the optimum center point for the number of clusters specified as k that group rows within a minimum distance for each cluster. Inclusion in a cluster for each row is defined by the features specified in the list of column parameters (before the semi-colon). The distance from the cluster center point is measured by the formula specified in measure.</p> <ul style="list-style-type: none"> ▪ col-or-expr1, col-or-expr2, ... col-or-exprN:: are individual columns or expressions used as features of the vectors that define clusters. This list of parameters must: <ul style="list-style-type: none"> ▪ Consist solely of columns with numeric values. ▪ Contain at least one column, but can have any number of columns. ▪ End with a semi-colon (;). This indicates the end of feature parameters and the beginning of the remaining, well-known parameters. ▪ k: is the number of clusters to create ▪ iterations: is the maximum number of iterations to perform to optimize clusters. ▪ measure: is the name of the formula to use to define membership for each row in a given cluster. Valid measures are: <ul style="list-style-type: none"> ▪ euclidean ▪ manhattan ▪ cosine

Aggregate	Window	Function	Description
			<ul style="list-style-type: none"> ▪ tanimoto ▪ squaredeclidean
	✓	analytics.kmeans_observations(Number col-or-expr1[,Number col-or-expr2,...Number col-or-exprN]; Integer k, Integer iterations, String measure)	<p>Returns the ID of the cluster that each row belongs to for the number of clusters specified as k. Inclusion in a cluster for each row is defined by the features specified in the list of column parameters (before the semi-colon). The distance from the cluster center point is measured by the formula specified in measure.</p> <ul style="list-style-type: none"> ▪ col-or-expr1, col-or-expr2, ... col-or-exprN;: are individual columns or expressions used as features of the vectors that define clusters. This list of parameters must: <ul style="list-style-type: none"> ▪ Consist solely of columns with numeric values. ▪ Contain at least one column, but can have any number of columns. ▪ End with a semi-colon (;). This indicates the end of feature parameters and the beginning of the remaining, well-known parameters. ▪ k: is the number of clusters to create ▪ iterations: is the maximum number of iterations to perform to optimize clusters ▪ measure: is the name of the formula to use to define membership for each row in a given cluster. Valid measures are: <ul style="list-style-type: none"> ▪ euclidean ▪ manhattan ▪ cosine ▪ tanimoto ▪ squaredeclidean
	✓	analytics.ordinal(String	Returns an index number for each unique

Aggregate	Window	Function	Description
		col-or-expr)	string value in the specified column or expression. Typically, assigning an ordinal number allows computations to work with string columns.
✓		skew(Number col-or-expr)	<p>Returns a measure of the asymmetry from a normal probability distribution of the values for this column or expression within the current partition or the current window in this partition. This uses the adjusted Fisher-Pearson standardized moment coefficient.</p> <p>Positive values tend to indicate that more values are found left of the mean, while negative values imply the reverse.</p> <p>This returns 0 if there are not at least three rows in the dataset. This function can also return NaN if the calculation is undefined or not a number.</p>

9.3.4.1.3 Window functions over RANGES

In RAQL **SELECT aggregationFunction(x) OVER(ORDER BY y) FROM s** is semantically equivalent to **SELECT aggregationFunction (x) OVER(ORDER BY y ROWS BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW) FROM s** . The SQL Standard however defines the semantics to correspond with **SELECT aggregationFunction (x) OVER(ORDER BY y RANGE BETWEEN UNBOUNDED PRECEDING AND CURRENT ROW) FROM s** which means that not only the current row but all "peer" rows (rows identical to the current row with respect to the ordering) are contained in the window. That means that a rolling aggregate without an explicit window frame specification currently yields different results in RAQL and any SQL standard compliant DBMS.

In order to ensure a high level of compliance with the SQL standard, window frames are also allowed to be specified using the **RANGE** key word.

Semantics

Window frames defined using the **ROWS** keyword are easy to understand. The window frame clause **ROWS BETWEEN x PRECEDING AND y FOLLOWING** simply defines two offsets **x** and **y** that for each row in the input partition define the first and the last row number of the corresponding window relative to the current row number. **PRECEDING** and **FOLLOWING** only

indicate whether the offset is negative (the row number precedes the current row number) or positive respectively.

Table 2. Example

rowNumber	ROWS BETWEEN 3 PRECEDING AND 3 FOLLOWING	ROWS BETWEEN 3 PRECEDING AND 1 PRECEDING	ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING
1	Window = [1,4]	Window = []	Window = [1,10]
2	Window = [1,5]	Window = [1,1]	Window = [2,10]
3	Window = [1,6]	Window = [1,2]	Window = [3,10]
4	Window = [1,7] (from row number 4-3 to 4+3)	Window = [1,3]	Window = [4,10]
5	Window = [2,8]	Window = [2,4]	Window = [5,10]
6	Window = [3,9]	Window = [3,5]	Window = [6,10]
7	Window = [4,10]	Window = [4,6]	Window = [7,10]
8	Window = [5,10]	Window = [5,7]	Window = [8,10]
9	Window = [6,10]	Window = [6,8]	Window = [9,10]
10	Window = [7,10]	Window = [7,9]	Window = [10,10]

On the other hand the semantics of window frame specifications using the **RANGE** keyword does not depend on the physical row numbers but on the actual values of a column. Still, the window frame clause defines two offsets **x** and **y**, but these offsets are not added to the row number but to an actual column value. Hence, any window frame specification involving **RANGE** and an offset other than **UNBOUNDED PRECEDING**, **UNBOUNDED FOLLOWING** or **CURRENT ROW** requires the window to be ordered on a single column and this column needs to have a numeric data type (in order to allow for offset arithmetics).

Table 3. Example

rowNumber	salary	SELECT avg(salary) OVER (ORDER BY salary RANGE BETWEEN 300 PRECEDING AND 300 FOLLOWING) FROM ...	considered window	
1	1750	1850	[1450,2050]	salaries between 1750 - 300 and 1750 + 300
2	1900	1850	[1600,2200]	
3	1900	1850	[1600,2200]	
4	2200	2100	[1900,2500]	

rowNumber	salary	SELECT avg(salary) OVER (ORDER BY salary RANGE BETWEEN 300 PRECEDING AND 300 FOLLOWING) FROM ...	considered window	
5	2400	2300	[2100,2700]	
6	2750	2866,67	[2450,3050]	
7	2900	2980	[2600,3200]	
8	2950	2980	[2650,3250]	
9	3100	2980	[2800,3400]	
10	3200	2980	[2900,3500]	

Table 4. Example

rowNumber	salary	SELECT avg(salary) OVER (ORDER BY salary RANGE BETWEEN 300 PRECEDING AND 300 FOLLOWING) FROM ...	considered window	
1	1750	1	[1750,1750]	salaries between 1750 - 0 and 1750 + 0
2	1900	2	[1900,1900]	
3	1900	2	[1900,1900]	
4	2200	1	[2200,2100]	
5	2400	1	[2400,2400]	
6	2750	1	[2750,2750]	
7	2900	1	[2900,2900]	
8	2950	1	[2950,2950]	
9	3100	1	[3100,3100]	
10	3200	1	[3200,3200]	

Note: The offsets in a **RANGE** frame clause define logical offsets in the ordered sequence defined by the order by clause. If the order by clause specifies a descending order (see next example) the offset **10 PRECEDING** translates to "a value that precedes the current value in the current sequence and differs at most by 10 (years/units/...)"

Table 5. Example

rowNumber	age	SELECT COUNT(*) OVER (ORDER BY age DESC RANGE 10 PRECEDING) FROM ...	considered window	
1	75	1	[85, 75]	ages between 75 + 10 and 75 - 0
2	68	2	[78, 68]	
3	68	2	[72, 62]	
4	56	2	[66, 56]	
5	55	3	[65, 55]	
6	51	3	[61, 51]	
7	48	5	[58, 48]	
8	48	5	[58, 48]	
9	47	6	[57, 47]	
10	43	5	[53, 43]	

Affected Queries

As mentioned above, the new **RANGE** specifier was introduced to ensure compliance with the SQL standard. There are however only very few cases where existing queries should be affected by those changes. Only queries specifying a window function with

- an explicit order-by specification and
- NO explicit window frame specification may be affected if
- The expressions defined in the order-by clause are not unique for each row.

If the order-by expressions are unique, than for each row there are no peers with respect to the ordering and the implicitly added window frame clause **RANGE UNBOUNDED PRECEDING** is equivalent to **ROWS UNBOUNDED PRECEDING** and hence the query results would be the same as before. If however the order-by specification is not unique then there may be rows having one or more peer rows so that the query result differs from the result in former RAQL releases and the user might want to explicitly add the window frame clause **ROWS UNBOUNDED PRECEDING** to yield the former results.

9.3.4.1.4 Migrate RAQL queries from Version 3.7 to 3.8

In version 3.8 significant changes to RAQL have been introduced that require changes to existing RAQL queries from 3.7 or earlier. These changes affect plain as well as analytic functions.

Note: In a few cases, the functions listed are supported but the parameters or the semantics of the function have changed sufficiently that queries may need updates to perform as expected.

9.3.4.1.4.1 Update plain functions for 3.8

Deprecated or Unsupported 3.7 Functions	Replace with 3.8 Function
concat(String column-or-literal, String column-or-literal)	Use the operator
"date"(String column) "date"(String column, String format)	to_date
"day"(Date column) "day"(String column) "day"(Long column)	extract_day
decimal(Double column) (Unsupported Function) decimal(Long column) decimal(String column)	cast(column as datatype)
"hour"(Date column) "hour"(String column) "hour"(Long column)	extract_hour
length(String column) (Unsupported Function)	char_length or character_length
"minute"(Date column) "minute"(String column) "minute"(Long column)	extract_minute
"month"(Date column) "month"(String column) "month"(Long column) "month"(Object column)	extract_month
number(String column) (Unsupported Function)	cast(column as datatype)
round(Object column)	This function no longer accepts objects and implicitly casts them to numeric values.

Deprecated or Unsupported 3.7 Functions	Replace with 3.8 Function
	<p>Mashups that relied on this implicit casting must update queries to use the following:</p> <pre>round(CAST(Object column AS double))</pre> <p>If no number of decimal places is specified, this function now rounds to zero decimals which may also affect mashup results. Mashups that require two decimal places must update queries to use the following:</p> <pre>round(Number column, "2")</pre>
<p>"second"(Date column)</p> <p>"second"(String column)</p> <p>"second"(Long column)</p>	extract_second
"string"(String column)	cast(column as datatype)
substr(String column, Integer begin, Integer end)	<p>substring(String column, Integer start-position, Integer length)</p> <p>Note: This is not a simple substitution as the semantics of the parameters that identify the characters to extract have different meanings in these functions.</p> <p>In addition to different parameters, character positions in substr use zero-based indexes while substring uses 1-based indexes.</p>
to_long(String column)	cast(column as datatype)
<p>"week"(Date column)</p> <p>"week"(String column)</p> <p>"week"(Long column)</p>	extract_week
<p>"year"(Date column)</p> <p>"year"(String column)</p> <p>"year"(Long column)</p> <p>"year"(Object column)</p>	extract_year

9.3.4.1.4.2 Update analytic functions for 3.8

Deprecated or Unsupported 3.7 Functions	Replace with 3.8 Function
correlation(Number column, Number column)	corr(Number column, Number column)

Deprecated or Unsupported 3.7 Functions	Replace with 3.8 Function
covariance(Number column, Number column)	covar(Number column, Number column)
denserank()	dense_rank()
firstvalue(Object column)	first_value(Object column)
analytics.kmean_clusters(String column-list, Integer k, Integer iterations, String measure)	analytics.kmeans_clusters(String column1[,String column2,...String columnN]; Integer k, Integer iterations, String measure) This changes the function name (for consistency) and changes the signature of the function. The syntax change supports a variable number of columns as parameters to identify the features for clustering.
analytics.kmeans_observations(String column-list, Integer k, Integer iterations, String measure)	analytics.kmeans_observations(String column1[,String column2,...String columnN]; Integer k, Integer iterations, String measure) This changes the signature of the function to support a variable number of columns as parameters to identify the features for clustering.
lastvalue(Object column)	last_value(Object column)
mean(Number column) (Unsupported Function)	avg(Number column)
rownumber()	row_number()

9.3.4.2 Creating new user-defined functions

You can define your own functions to use in RAQL queries in addition to the RAQL function extensions (page 404). User-defined functions are Java classes that you write, configure and deploy to MashZone NextGen.

User-defined function can be:

- **Plain functions:** used in Select, Over, Where, Order By and Group By clauses in RAQL queries. They typically either cast (change) the datatype of column values, extract part of the values or transform values in some way.
Plain functions are applied individually to each value in the column specified without access to values in other rows.
- **Window analytic functions:** use some of the rows in a partition or window to perform a calculation and return a value for each row. Typically, this uses rows that are relative to the current row, such as **first_value** or **row_number**.

You must include an `Over` clause in queries that use window analytic functions. This defines the partitions or windows used by the function.

- **Aggregate analytic functions:** use all rows in the dataset, group, partition or window to perform a calculation and return a single value. For example, `sum` adds the values of all rows in the current scope.

Aggregate analytic functions can be used in the `Select` or `Having` clauses of RAQL queries:

- To return a single value for each group defined in a `Group By` clause.
- To return a single value for each partition or each window defined in an `Over` clause. This single value is added to each row in the partition or window. You can also use aggregate analytic functions to return running calculations for partitions or windows defined in an **Over** clause.
- To return a single value for the entire dataset if no group definition or partition definition is specified.

To write user-defined functions for RAQL, you should [Set Up Your Development Environment](#) (page 420).

Then:

1. [Write Plain Functions for RAQL](#) (page 421)
2. [Write Window Analytic Functions for RAQL](#) (page 425)
3. [Write Aggregate Analytic Functions for RAQL](#) (page 432)
4. [Use a Factory for Function Overloading](#) (page 437), if needed
5. [Configure, Compile, Deploy and Test User-Defined Functions](#) (page 422)

9.3.4.2.1 Set Up Your Development Environment

To get started, you need:

- The latest Java Development Kit 8. See [JDK 1.8](#) to download and install this, if needed.
- A folder for your user-defined function **library** with this structure:

```
MyOrgRaqlLib
lib
src
  com
    MyOrg
      raqlUdf
lib.json
```

- User-defined functions are packaged and deployed in named libraries. The library name also uniquely defines your user-defined functions from built-in functions or user-defined functions in other libraries that have been deployed in MashZone NextGen. Library names must match the name of the library folder containing your source code.

- Valid library names must be unique for a MashZone NextGen Server. They must start with a letter and can contain letters, numbers or underscores (_).
Important: The name `analytics` is **reserved** for the MashZone NextGen UDF function library.
- You can organize user-defined functions however you need. User-defined functions can be packaged in multiple libraries. Each library can contain one or more Java packages. Each package can contain multiple classes. For plain functions, each class can contain one or more functions. For analytic functions, each function is packaged as a class.
Note: You **cannot** use periods (.) in either library or user-defined function names as this character is the reserved separator between library and function names.
- JAR files for any third-party Java libraries that your user-defined functions depend on, placed in the **library-name/lib** folder.
- Update your classpath to include the following JAR files:
 - **web-apps-home/mashzone/WEB-INF/lib/jackbe-presto-raql-*version*.jar** for RAQL,
 - **web-apps-home/mashzone/WEB-INF/lib/rtm-core-*version*.jar** for the UDF interfaces,
 - and any third party JAR files you add to **library-name/lib**.
- For an example, see [Configure, Compile, Deploy and Test User-Defined Functions](#) (page 422).

Important: It is **not** a good practice to copy either the RAQL or UDF interface JAR files to your **library-name /lib** folder to simplify the classpath. This can cause errors when you deploy your user-defined function library. See [Configure, Compile, Deploy and Test User-Defined Functions](#) (page 422) for more information.

Before you compile, deploy and test your user-defined functions, you will also need a library configuration file (`lib.json`). We will cover these requirements later, in [Configure, Compile, Deploy and Test User-Defined Functions](#) (page 422).

9.3.4.2.2 Write Plain Functions for RAQL

Effective in version 3.8, plain functions implement the `UserDefinedFunctionAdapter` interface in the `de.rtm.push.adapters` package of the MashZone NextGen RAQL User Defined Function API. If functions need to support multiple signatures, you can also implement the `UserDefinedFunctionAdapterFactory` interface.

Note: These interfaces are not backwards compatible with the UDF API from MashZone NextGen 3.7 or earlier.

You can implement plain functions more simply as Java classes and let RAQL automatically derive the methods required by the interface. Thus at a minimum plain functions must:

- Import the `com.jackbe.jbp.raql.udx.loader.RaqlFunc` annotation interface.
- Include annotations to identify which methods are RAQL functions.
- Be implemented as public static class methods that return primitive types.

This example contains two plain functions, `replace` and `capitalize`:

```
package com.jackbe.jbp.raql.samplelib.annotated.udf;
```

```
/*
 * Sample class in a library of plain user-defined functions for RAQL
 */

import com.jackbe.jbp.raql.udx.loader.RaqlFunc;
public class StringFunctions {

    @RaqlFunc
    public static String replace(String val, String oldStr, String newStr) {
        if (val == null || oldStr == null || newStr == null)
            return null;

        return val.replace(oldStr, newStr);
    }

    @RaqlFunc(name="capitalize")
    public static String upper(String val) {
        return val == null ? null : val.toUpperCase();
    }
}
```

The `@RaqlFunc` annotation identifies which methods should be associated with RAQL functions. If you omit the name parameter, the name of the method becomes the name of the RAQL function within this UDF library. Use (name="**alias**") to use a different name for the function from the method name.

Also, plain user-defined functions may be annotated using the `@RaqlFunc` annotation. If within one single class multiple static methods with distinct signatures are annotated with the `@RaqlFunc` annotation and are given the same name, they will all be registered under that name and the engine will choose the most suitable signature for a particular call.

Tip: Using annotations to configure the methods for user-defined functions in RAQL is a best practice. However, you can skip the annotations in your Java classes and instead provide configuration that maps methods to user-defined functions in the `lib.json` configuration file for your UDF libraries. See External UDF Library Deployment Folder (page 425) for more information.

With plan functions, you can annotate several static methods with different signatures in a single class using the same function name. All of the methods are registered under that single function name. At runtime, the method with the appropriate signature is used.

9.3.4.2.3 Configure, Compile, Deploy and Test User-Defined Functions

To compile and deploy user-defined functions

1. Add a folder with the name of your library under the default external UDF library deployment folder **<MashZone NextGen installation>/raql-udfs**.

Note: The folder for UDF library deployment is set in the system property `-Dpresto.raql.udf.libsDir`. This may be a different folder in clustered environments with a shared external configuration folder for MashZone NextGen.

This new folder is the root deployment folder for your library. For example:

/SoftwareAG/MashZoneNG/raql-udfs/MyOrgRaqlLib

2. Complete configuration that identifies the Java packages for this library in a lib.json file. This file must reside in the root deployment folder for your library, created in the previous step.

Note: You may also want to place a copy of this configuration in your library development folder to track in your source control system.

This file uses the JSON format to identify packages that contain user-defined functions for a library. For example:

```
{
  "exportedFuncs" : {
    "annotatedPackages" : [
      "com.jackbe.jbp.raql.samplelib.annotated.udf",
      "com.jackbe.jbp.raql.samplelib.annotated.uda",
      "com.jackbe.jbp.raql.samplelib.annotated.udw"
    ],
    "classes" : [
      {
        "name" : "com.jackbe.jbp.raql.samplelib.configured
          .MoreStringFunctions",
        "funcs" : [
          { "name" : "low", "method" : "lower" },
          { "method" : "hashCode" }
        ]
      },
      {
        "name" : "com.jackbe.jbp.raql.samplelib.configured.
          AverageAggregationFunctionFactory",
        "funcs" : [
          { "name" : "avg" }
        ]
      },
      {
        "name" : "com.jackbe.jbp.raql.samplelib.configured.
          LeadWindowFunction",
        "funcs" : [
          { "name" : "lead" }
        ]
      }
    ]
  }
}
```

This example identifies three packages for this library. The classes property must be present in configuration, but can be an empty array. The RAQL Engine will search each class within the identified packages for function annotations to find the functions to add.

You can also optionally identify classes, functions and methods in this configuration. See UDF Library Configuration (page 424) for more information on this method.

3. Compile your Java class with user-defined functions, being sure to include the jackbe-presto-raql-**version**.jar and rtm-core-**version**.jar files and any third-party libraries you used in this class in the classpath.

For **Windows** for example:

```
%JAVA_HOME%/bin/javac -classpath
c:\SoftwareAG\MashZoneNG\apache-tomcat\webapps\mashzone\WEB-INF\lib\jackbe-presto-
to-raql-10.3.jar;c:\SoftwareAG\MashZoneNG\apache-tomcat\webapps\mashzone\WEB-INF\lib\rtm-core-10.3.jar -d classes src\com\MyOrg\raqlUdf\MyReplaceFunc.java
```

For **Linux, OS/X or UNIX** for example:

```
javac -classpath
/SoftwareAG/MashZoneNG/apache-tomcat/webapps/mashzone/WEB-INF/lib/jackbe-presto-
-raql-10.3.jar,/SoftwareAG/MashZoneNG/apache-tomcat/webapps/mashzone/WEB-INF/lib
/rtm-core-10.3.jar -d classes src/com/MyOrg/raqlUdf/MyStringFuncs.java
```

4. This will add the compiled class to the classes folder in your development folder for this library.

5. Deploy the classes and any third-party libraries for this library to the MashZone NextGen Server. Copy the following folders to the root deployment folder for your library:

/classes folder and all contents

/lib folder and all contents

Important: If you have included either the RAQL or RTM JAR files in your **library-name** /lib folder, do not copy these JARs to the deployment folder for your library as this causes errors when functions in the library are used.

6. Restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) (page 389 → link to correct chapter) for instructions.
7. Write RAQL queries that use these new user-defined functions to test them.
You refer to user-defined functions in the form **library-name.function-name([arg,...])** to identify both the library name and the function name.

9.3.4.2.3.1 UDF Library Configuration

The lib.json file contains configuration that identifies the Java packages with user-defined functions for RAQL. It can also optionally contain configuration that identifies the specific classes and methods in these packages and the function names to map to methods.

Note: It is a best practice to configure methods and function names using annotations in the Java classes directly. If you choose to use lib.json configuration instead, you should omit the annotations in your Java classes for user-defined functions.

The following example includes method and function mapping configuration along with the required package configuration information:

```
{
  "exportedFuncs": {
    "annotatedPackages": [
      "com.MyOrg.raqlUdf",
      "com.MyOrg.aggregate.raqlUdf",
      "com.MyOrg.window.raqlUdf" ],

```

```
"classes": [
  { "name": "com.MyOrg.raqlUDF.MyStringFuncs",
    "funcs": [ { "method": "replace" },
      { "name": "capitalize", "method": "upper" } ] },
  { "name": "com.MyOrg.aggRaqlUdf.KurtosisFunction",
    "funcs": [ { "name": "kurtosis" } ] }
]
}
}
```

9.3.4.2.3.2 External UDF Library Deployment Folder

A default external UDF Library Deployment folder is created when you install MashZone NextGen at **<MashZone NextGen installation>/raql-udfs** where you can deploy all your user-defined functions.

In clustered environments, you may want to create a shared external folder for MashZone NextGen configuration and **move** all user-defined functions, including the MashZone NextGen built-in function library, to this shared location for all members of the cluster.

If you move user-defined functions from the default UDF Library Deployment folder, you must also update an environmental variable for each MashZone NextGen Server:

1. Edit the script for the appropriate operating system in any text editor of your choice:
<MashZone NextGen installation>/apache-tomcat/bin/setenv.bat file, for Windows systems, or
<MashZone NextGen installation>/apache-tomcat/bin/setenv.sh file for Linux, OS/X or UNIX systems.
2. Add or update the `-Dpresto.raql.udf.libsDir` system property with the path to point to the new shared location.
3. Save your changes and restart the MashZone NextGen Server. See Start and Stop the MashZone NextGen Server (page 113) for instructions.

9.3.4.2.4 Write Window Analytic Functions for RAQL

With window analytic functions, each function is a single class that implements the `UserDefinedWindowFunctionAdapter` interface (`de.rtm.push.adapters.windowfunctions` package) from the MashZone NextGen RAQL User Defined Function API.

Window analytic functions, unlike plain functions, have access to all rows, or **records** within the current partition or window which they can use to perform calculations. Unlike aggregate analytic functions, however, they provide a different calculation for each record.

To accomplish this, window analytic functions use the following methods:

- **createInitialState()**: to reset state for the current window or partition.
- **checkWindowSpecification(boolean isPartitioned, boolean isOrdered, WindowFrameSpecification windowFrameSpec)**: to validate that the window definition meets requirements such as being sorted.
- **call(S state, UserDefinedWindowFunctionAdapter.PartitionEntry currentEntry, UserDefinedWindowFunctionAdapter.Partition partition, int currentIndex, WindowFrame currentWindowSpec, WindowFrame prevWindowSpec)**: window analytic functions must implement this method with the core logic of the function and return the result of the calculation as **WindowFunctionResult<S,O>**.

There are also basic 'housekeeping' methods: **getParameterTypes()** and **getReturnType()**.

You set up your development environment for window analytic functions just the same as for plain functions. See [Set Up Your Development Environment](#) (page 420) for details.

We're going to use two examples. The first example shows the basics of a window analytics function and how to track state and set the function result. This example implements a simple sum. To create this function, you:

1. Construct and Initialize the Window Analytic Function Class (page 426)
2. Implement the call Method for the Window Analytic Function (page 427)

And Configure, Compile, Deploy and Test User-Defined Functions (page 422). For the complete code, see [Complete WindowSum Example](#) (page 429).

The second example implements the MashZone NextGen built-in lead function which illustrates techniques to [Work with Specific Records in Window Calculations](#) (page 430) using the current position of a record.

9.3.4.2.4.1 Construct and Initialize the Window Analytic Function Class

This example, `MySumWindowFunction.java` is available in the sample user-defined functions package at [<MashZone NextGen installation>/raql-udfs/SampleRaqlLib](#).

Your window analytic function class imports the RAQL UDF annotation class, `com.jackbe.jbp.raql.udx.loader.RaqlFunc`, various classes in the MashZone NextGen RAQL User Defined Function API and implements the `UserDefinedWindowFunctionAdapter` interface:

```
package com.raql.samples;

import com.jackbe.jbp.raql.udx.loader.RaqlFunc;

import de.rtm.push.adapters.Adapters;
import de.rtm.push.adapters.windowfunctions.UserDefinedWindowFunctionAdapter;
import de.rtm.push.adapters.windowfunctions.WindowFrame;
import de.rtm.push.adapters.windowfunctions.WindowFrameSpecification;
import de.rtm.push.adapters.windowfunctions.WindowFunctionResult;
import de.rtm.util.exception.IncompatibleWindowSpecificationException;

/**
 * This window function adapter computes the sum of a window.
 */
```

```
@RaqlFunc(name="mySumFunction")
public class WindowSum implements UserDefinedWindowFunctionAdapter
<Double, Double> {

    protected final Type[] parameterTypes;

    public WindowSum(Type[] parameterTypes) {
        this.parameterTypes = parameterTypes;
    }

    /**
     * Returns the type of the result, which is simply always Double.
     */
    @Override public Type getReturnType() {
        return Type.DOUBLE;
    }

    /**
     * Returns the types of the input parameters.
     */
    @Override public Type[] getParameterTypes() {
        return parameterTypes;
    }

    /**
     * Creates the initial state of the window sum, which is 0.
     */
    @Override
    public Double createInitialState() {
        return 0d;
    }
}
```

This example overrides the default implementation for createInitialState to reset the sum to zero.

9.3.4.2.4.2 Implement the call Method for the Window Analytic Function

You implement the call method with the core logic for your window analytic function. The state parameter represents either the initial state (for the first window evaluation) or the state computed from the previous window evaluation.

Note: User-defined functions should be stateless wherever possible as a best practice. Instead, you can manage intermediate state information for window calculations in the state property of the WindowFunctionResult.

It also has parameters for the current row within the current partition, the current partition, the index for the current row and specifications for the current window and the previous window that define the context for the function.

```
...
/**
 * Computes the sum for the current window. Instead of simply summing
 * up the values of the current window, only the required values
 * from current and previous window are combined with the state of
 * the previous window evaluation. This approach allows for a more
 * efficient evaluation.
 *
 * @param state The state which is either the initial state for
```

```

*           the first window evaluation or the state as
*           computed during the previous window evaluation
* @param currentEntry The current row of the partition being
*           processed
* @param partition The partition being processed
* @param currentIndex The index of the current row being
*           processed
* @param currentWindowSpec The specification of the current window
*           frame within the partition
* @param previousWindowSpec The specification of the previous window
*           frame within the partition
* @return The result of the window function for the current window
*/
@Override
public WindowFunctionResult<Double, Double> call(Double state,
        PartitionEntry currentEntry, Partition partition, int currentIndex,
        WindowFrame currentWindowSpec, WindowFrame previousWindowSpec) {
    // initialize new sum with old sum from state
    Double newSum = state;
    // if previous window has values, remove any from the subtotal that
    // are not in current window
    boolean valuesInPreviousWindow = previousWindowSpec != null &&
        !previousWindowSpec.isEmpty();
    if (valuesInPreviousWindow) {
        final int removeFrom = Math.min(currentWindowSpec.getStartIndex(),
            previousWindowSpec.getEndIndex());
        for (int i = previousWindowSpec.getStartIndex(); i < removeFrom; i++) {
            newSum -= ((Number) partition.get(i).getColumnValue(0)).doubleValue();
        }
    }
    // and add values from current window that are not in previous window
    final int addFrom = !valuesInPreviousWindow?currentWindowSpec.getStartIndex()
        :Math.max(currentWindowSpec.getStartIndex(),
previousWindowSpec.getEndIndex());
    for (int i = addFrom; i < currentWindowSpec.getEndIndex(); i++) {
        newSum += ((Number)partition.get(i).getColumnValue(0)).doubleValue();
    }
    // return the new sum as state for use in the next window
    // evaluation and as result for the current row
    return Adapters.createWindowFunctionResult(newSum, newSum);
}

@Override
public void checkWindowSpecification(boolean isPartitioned, boolean isOrdered,
        WindowFrameSpecification windowFrameSpecification) throws
        IncompatibleWindowSpecificationException {
    // nothing to be done
}
}

```

The call method returns an object that implements the WindowFunctionResult interface.

WindowFunctionResult instances contain:

- The result of the function, in this case the sum of this numeric column for the current row in the current window and partition.
- State information needed to apply this function to the next row.
In this example, the state is also simply the result of the function. The calculation of the sum for each row uses this as a starting point and then backs out values for any rows that are no longer considered part of the window and adds in values for any new rows in the window.

Lastly, this implements the `checkWindowSpecification` method that is required for the interface as a no op method.

9.3.4.2.4.3 Complete WindowSum Example

```
package com.raql.samples;

import com.jackbe.jbp.raql.udx.loader.RaqlFunc;
import de.rtm.push.adapters.Adapters;
import de.rtm.push.adapters.windowfunctions.UserDefinedWindowFunctionAdapter;
import de.rtm.push.adapters.windowfunctions.WindowFrame;
import de.rtm.push.adapters.windowfunctions.WindowFrameSpecification;
import de.rtm.push.adapters.windowfunctions.WindowFunctionResult;
import de.rtm.util.exception.IncompatibleWindowSpecificationException;

/**
 * This window function adapter computes the sum of a window.
 */
@RaqlFunc(name="mySumFunction")
public class WindowSum implements UserDefinedWindowFunctionAdapter<Double, Double>
{
    protected final Type[] parameterTypes;

    public WindowSum(Type[] parameterTypes) {
        this.parameterTypes = parameterTypes;
    }

    /**
     * Returns the type of the result, which is simply always Double.
     */
    @Override public Type getReturnType() {
        return Type.DOUBLE;
    }

    /**
     * Returns the types of the input parameters.
     */
    @Override public Type[] getParameterTypes() {
        return parameterTypes;
    }

    /**
     * Creates the initial state of the window sum, which is 0.
     */
    @Override
    public Double createInitialState() {
        return 0d;
    }

    /**
     * Computes the sum for the current window. Instead of simply summing
     * up the values of the current window, only the required values
     * from current and previous window are combined with the state of
     * the previous window evaluation. This approach allows for a more
     * efficient evaluation.
     *
     * @param state The state which is either the initial state for
     * the first window evaluation or the state as
     * computed during the previous window evaluation
     * @param currentEntry The current row of the partition being
```

```

*           processed
* @param partition The partition being processed
* @param currentIndex The index of the current row being
*           processed
* @param currentWindowSpec The specification of the current window
*           frame within the partition
* @param previousWindowSpec The specification of the previous window
*           frame within the partition
* @return The result of the window function for the current window
*/
@Override
public WindowFunctionResult<Double, Double> call(Double state,
        PartitionEntry currentEntry, Partition partition, int currentIndex,
        WindowFrame currentWindowSpec, WindowFrame previousWindowSpec) {
    // initialize new sum with old sum from state
    Double newSum = state;
    // if previous window has values, remove any from the subtotal that
    // are not in current window
    boolean valuesInPreviousWindow = previousWindowSpec != null &&
        !previousWindowSpec.isEmpty();
    if (valuesInPreviousWindow) {
        final int removeTo = Math.min(currentWindowSpec.getStartIndex(),
            previousWindowSpec.getEndIndex());
        for (int i = previousWindowSpec.getStartIndex(); i < removeTo; i++) {
            newSum -= ((Number) partition.get(i).getColumnValue(0)).doubleValue();
        }
    }
    // and add values from current window that are not in previous window
    final int addFrom = !valuesInPreviousWindow?currentWindowSpec.getStartIndex()
        :Math.max(currentWindowSpec.getStartIndex(),
previousWindowSpec.getEndIndex());
    for (int i = addFrom; i < currentWindowSpec.getEndIndex(); i++) {
        newSum += ((Number)partition.get(i).getColumnValue(0)).doubleValue();
    }
    // return the new sum as state for use in the next window
    // evaluation and as result for the current row
    return Adapters.createWindowFunctionResult(newSum, newSum);
}
@Override
public void checkWindowSpecification(boolean isPartitioned, boolean isOrdered,
WindowFrameSpecification windowFrameSpecification) throws
    IncompatibleWindowSpecificationException {
    // nothing to be done
}
}

```

9.3.4.2.4.4 Work with Specific Records in Window Calculations

This example is a window analytics function similar to the MashZone NextGen built-in lead function which returns the value for the specified column for a row (a record) that follows the current row by a specific offset:

```

package com.raql.samples;

import java.io.Serializable;
import com.jackbe.jbp.raql.udx.loader.RaqlFunc;
import de.rtm.push.adapters.Adapters;

```

```
import de.rtm.push.adapters.windowfunctions.UserDefinedWindowFunctionAdapter;
import de.rtm.push.adapters.windowfunctions.WindowFrame;
import de.rtm.push.adapters.windowfunctions.WindowFrameSpecification;
import de.rtm.push.adapters.windowfunctions.WindowFunctionResult;
import de.rtm.util.exception.IncompatibleWindowSpecificationException;

/**
 * This window function adapter determines the lead value for a
 * specific row based on an offset within a window.
 */
@RaqlFunc(name="myLeadFunction")
public class Lead implements UserDefinedWindowFunctionAdapter
    <Serializable, Serializable> {
    private final static Serializable DEFAULT_DEFAULT_VALUE = null;
    private final static int DEFAULT_OFFSET = 1;

    protected final Type returnType;
    protected final Type[] parameterTypes;

    public Lead(Type[] parameterTypes) {
        this.returnType = parameterTypes[0];
        this.parameterTypes = parameterTypes;
    }

    /**
     * Returns the type of the result.
     */
    @Override
    public Type getReturnType() {
        return returnType;
    }

    /**
     * Returns the types of the input parameters.
     */
    @Override
    public Type[] getParameterTypes() {
        return parameterTypes;
    }

    /**
     * Creates the initial state.
     */
    @Override
    public Serializable createInitialState() {
        return null;
    }

    /**
     * Determines the lead value of the current window. This has two
     * optional parameters: the offset and the default value ?where?.
     *
     * @return The result of the window function for the current window
     */
    @Override
    public WindowFunctionResult<Serializable, Serializable> call(Serializable state,
        PartitionEntry currentEntry, Partition partition, int currentIndex,
        WindowFrame currentWindowSpec, WindowFrame previousWindowSpec) {
        final Serializable[] columnValues = currentEntry.getColumnValues();
        // when the offset is not specified, use the default offset, which is 1
        final int offset = columnValues.length > 1 ? (Integer)columnValues[1] :
            DEFAULT_OFFSET;
        final int index = currentIndex + offset;
    }
}
```

```

final Serializable result;
if (index < currentWindowSpec.getStartIndex() || index >=
    currentWindowSpec.getEndIndex())
    // when the default value is not specified, use the default default value,
    // which is null
    result = columnValues.length > 2 ? columnValues[2] : DEFAULT_DEFAULT_VALUE;
else
    result = partition.get(index).getColumnValue(0);
return Adapters.createWindowFunctionResult(result, result);
}
@Override
public void checkWindowSpecification(boolean isPartitioned, boolean isOrdered,
    WindowFrameSpecification windowFrameSpecification) throws
    IncompatibleWindowSpecificationException {
    // nothing to be done
}
}

```

To work with a specific record relative to the current record, this function uses both the index of the current record as well as the specification of the current window.

9.3.4.2.5 Write Aggregate Analytic Functions for RAQL

Aggregate analytic functions are a single class that implements the **UserDefinedAggregationFunctionAdapter** interface in the `de.rtm.push.adapters` package of the MashZone NextGen RAQL User Defined Function API.

Like window analytic functions, aggregate analytic functions have access to all rows, or **records** in the current group, partition or window. They perform an aggregate calculation that uses all records to return a single result. Depending on where they are used in RAQL queries, this single result may be all the query returns or it may be returned as a column on every row.

To accomplish this, aggregate analytic functions use the following methods:

- **createInitialState()**: is optional. This resets the aggregate state for the current window, partition or group to whatever value or state the calculation should begin from.
- **aggregate(S state, Serializable[] values)**: performs an intermediate calculation using the current state and the value of the next record to obtain a new intermediate state.
- **getAggregate(S state)**: implements the final computation for the group, partition or window, if any, and returns the final aggregate result once all intermediate calculations are complete.

There are also several 'housekeeping' methods you must implement: **getAggregateType()**, **getParameterTypes()**, **isEmptyAggregate(S state)**, **negativeCall(S state, Serializable[] values)** and **supportsPN()**.

You set up your development environment for aggregate analytic functions just the same as for plain functions. See [Set Up Your Development Environment](#) (page 420) for details. Write the class for your aggregate analytic function and then [Configure, Compile, Deploy and Test User-Defined Functions](#) (page 422).

We're going to use two examples. The first example, [My Average Aggregate Example](#) (page 433) is an implementation of the built-in `avg(Number column)` method. This example shows the basics

of an aggregate analytics function, how to track state for intermediate steps and then perform the final calculation.

The second example, Kurtosis Using a Third Party Library (page 435), uses methods in the Apache Commons Math library to calculate the kurtosis for a column. This is an example of how to use third-party libraries.

9.3.4.2.5.1 My Average Aggregate Example

This example, **MyAverageAggregationFunction** is available in the sample user-defined functions package at **<MashZone NextGen installation>/raql-udfs/SampleRaqlLib**. It implements an aggregate function `myAverageFunction(Number column)` similar to the MashZone NextGen built-in average function:

The aggregate method performs intermediate calculations. This increments the number of rows processed and adds the current column value to a subtotal. Both of these intermediate values are added to the state object. The `getAggregate` method then uses state to calculate the group/partition/window average.

```
package com.raql.samples;

import java.io.Serializable;
import com.jackbe.jbp.raql.udx.loader.RaqlFunc;
import de.rtm.push.adapters.UserDefinedAggregationFunctionAdapter;

/**
 * This aggregate adapter computes the average of a group, partition
 * or window.
 */
@RaqlFunc(name="myAverageFunction")
public class MyAverageAdapter implements UserDefinedAggregationFunctionAdapter
<Number[]> {
    /**
     * Returns the types of the input parameters.
     */
    @Override
    public Type[] getParameterTypes() {
        return new Type[] {Type.DOUBLE};
    }
    /**
     * Returns the type of the result.
     */
    @Override
    public Type getAggregateType() {
        return Type.DOUBLE;
    }
    /**
     * Creates the initial state of the average.
     */
    @Override
    public Number[] createInitialState() {
        // the first value is the count and the second the cumulative sum
        return new Number[] {0l, 0d};
    }
    /**
     * Aggregates the current internal state with the new input values
```

```
* and derives a new internal state.
* @param state The current internal state
* @param values The new input values
* @return The new internal state
*/
@Override
public Number[] aggregate(Number[] state, Serializable[] values) {
    Long n = (Long) state[0];
    Double sum = (Double) state[1];
    // null-aware handling of input values
    if (values[0] != null) {
        sum += (Double) values[0];
        n += 1;
    }
    // always return a new state
    return new Number[] {n, sum};
}
/**
 * Derives the final aggregate value from the current internal state.
 * @param state the internal state
 * @return the aggregate value
 */
@Override
public Object getAggregate(Number[] state) {
    long n = (Long) state[0];
    double sum = (Double) state[1];
    // the average of an empty data set is null
    if (n == 0)
        return null;

    return sum/n;
}
/**
 * Removes the input values from the internal state.
 * This step is required for the support of the
 * positive/negative approach.
 * @param state the current internal state
 * @param values the input values
 * @return the new internal state
 */
@Override
public Number[] negativeCall(Number[] state, Serializable[] values) {
    if (values[0] == null)
        return state;

    long n = (Long) state[0];
    double sum = (Double) state[1];

    n--;
    sum -= (Double) values[0];

    // always return a new state
    return new Number[] {n, sum};
}
/**
 * Indicates support of positive/negative approach, which can add and
 * remove values from the internal state; aggregate analytic functions
 * supporting that approach allow for a more efficient evaluation
 * when windows are used.
 * @return indicates whether PN approach is supported

```

```
    */
    @Override
    public boolean supportsPN() {
        return true;
    }
    /**
     * Indicates whether the internal state is empty.
     * @param state the internal state
     * @return indicates whether the internal state is empty
     */
    @Override
    public boolean isEmptyAggregate(Number[] state) {
        if ((Long) state[0] == 0)
            return true;
        return false;
    }
}
```

This example illustrates that the state object can track multiple properties. In this case state tracks both the count of the number of rows processed so far and the subtotal of the values for the column being averaged.

9.3.4.2.5.2 Kurtosis Using a Third Party Library

Kurtosis is a statistical measure of 'peakedness' in the values for a dataset compared to a normal distribution. This indicates how closely the distribution matches the rounded bell shape of a normal distribution.

In this example, we will use an implementation of kurtosis provided in the Apache Commons Math library, version 2.2. The method to calculate kurtosis, in the DescriptiveStatistics class in the Apache Library, expects the values to use as the probability distribution to be primitive values in an array.

To support this, the aggregate method builds an array from the column values for records in a group, partition or window. The getAggregate method then uses this array to perform the calculation. As always, the state object is used to hold state for both methods.

```
package com.raql.samples;

import java.io.Serializable;
import org.apache.commons.math.stat.descriptive.DescriptiveStatistics;
import com.jackbe.jbp.raql.udx.loader.RaqlFunc;
import de.rtm.push.adapters.UserDefinedAggregationFunctionAdapter;

@RaqlFunc(name="myKurtosisFunction")
/**
 * This adapter computes the kurtosis for a group, partition or window.
 */
public class KurtosisAdapter implements
UserDefinedAggregationFunctionAdapter<double[]> {
    /**
     * Returns the types of the input parameters.
     */
    @Override
    public Type[] getParameterTypes() {
        return new Type[] {Type.DOUBLE};
    }
}
```

```
}
/**
 * Returns the type of the result.
 */
@Override
public Type getAggregateType() {
    return Type.DOUBLE;
}
/**
 * Creates the initial state of the kurtosis. Note that
 * the state solely stores the input values. Therefore
 * the initial state is an empty array.
 * @return the initial state of the kurtosis
 */
@Override
public double[] createInitialState() {
    return new double[] {};
}
/**
 * Aggregates the current internal state with the new input values
 * and derives a new internal state. As the kurtosis aggregate is
 * computed by an external library, the state solely stores
 * the incoming values.
 * @param state the current internal state
 * @param values the new input values
 * @return the new internal state
 */
@Override
public double[] aggregate(double[] state, Serializable[] values) {
    if (values[0] != null) {
        double[] newState = new double[state.length+1];
        System.arraycopy(state, 0, newState, 0, state.length);
        newState[newState.length-1] = (Double) values[0];
        return newState;
    }
    else
        return state;
}
/**
 * Derives the return value from the current internal state.
 * The kurtosis aggregate is computed by calling an external
 * library with all input values.
 * @param state the internal state
 * @return the kurtosis value
 */
@Override
public Object getAggregate(double[] state) {
    DescriptiveStatistics ds = new DescriptiveStatistics(state);
    double kurtosis = ds.getKurtosis();
    if (Double.isNaN(kurtosis))
        return 0.0;
    else
        return kurtosis;
}
/**
 * Removes the input values from the internal state.
 * This step is required for the support of the
 * positive/negative approach.
 * @param state the current internal state
 * @param values the input values
 */
```

```

    * @return the new internal state
    */
    @Override
    public double[] negativeCall(double[] state, Serializable[] values) {
        if (values[0] == null)
            return state;
        else {
            boolean gotValue = false;
            double[] newState = new double[state.length-1];
            double value = (Double)values[0];
            for (int i = 0; i < state.length; i++) {
                gotValue |= state[i] == value;
                newState[i] = state[i + (gotValue ? 1 : 0)];
            }
            return newState;
        }
    }

    /**
     * Indicates support for positive/negative approach, which allows
     * adding or removing values from the internal state;
     * This approach provides a more efficient evaluation
     * when windows are used.
     * @return indicates whether PN approach is supported
     */
    @Override
    public boolean supportsPN() {
        return true;
    }

    /**
     * Indicates whether the internal state is empty.
     * @param state the internal state
     * @return indicates whether the internal state is empty
     */
    @Override
    public boolean isEmptyAggregate(double[] state) {
        return state.length == 0;
    }
}

```

To compile this example, you must include the Apache Commons Math library, version 2.2, in the classpath. You may include the jar file for this library in the lib folder for the user-defined function library. This specific library is also used in MashZone NextGen, so you also simply add the jar file for this library to the classpath.

9.3.4.2.6 Use a Factory for Function Overloading

If a user-defined function must support multiple signatures, you must implement the corresponding factory interface:

- **UserDefinedFunctionAdapterFactory** for plain functions
- **UserDefinedAggregationFunctionAdapterFactory** for aggregate analytic functions
- **UserDefinedWindowFunctionAdapterFactory** for window analytic functions

For more information on the advantages of user-defined function factories and examples. See the RAQL User-Defined Function API reference.

9.4 MashZone NextGen Platform API Console

This console allows MashZone NextGen administrators and developers to discover and run platform APIs for MashZone NextGen. Platform APIs can be useful to investigate issues that you may report to Technical Support. They may also be used in some cases during upgrades to a new version of MashZone NextGen. Platform APIs can also be used to provide specialized features or integration.

Platform APIs are organized in services that focus on specific aspects of MashZone NextGen. Find and select the MashZone NextGen Platform service and method you want to run. In most cases, the service names are self-explanatory. See MashZone NextGen Platform Services (page 438) for more information if needed.

Once you select a platform API method, a sample request to invoke this method fills the Request pane.

9.4.1 Open the MashZone NextGen API Console

MashZone NextGen includes the API Console (page 438), a web based tool for MashZone NextGen administrators to discover and run platform APIs.

Procedure

1. Click the  user icon in the program bar.
2. Click **API Console**.

The API Console opens.

9.4.2 MashZone NextGen Platform Services

Service	Description
ConfigService	Contains methods to get and set configuration properties for the MashZone NextGen Server. Includes both methods for specific properties and generic methods to get and set properties by name.
DashboardFeedService	Contains methods to search and to delete dashboards and data feeds. The searchDashboards and searchFeeds methods deliver a list of dashboards which fits to the specified search criteria. The deleteDashboards and deleteFeeds methods delete all dashboards or data feeds which fits to specific GUIDs. Search parameter: <ul style="list-style-type: none"> ▪ keyword - The search term. If the parameter is empty, all dashboards matching the other criteria are returned, otherwise, only dashboards that additionally match the search term are returned.

Service	Description
	<ul style="list-style-type: none"> ▪ timestamp - If this parameter is > 0, only dashboards are listed that have been changed after the specified timestamp. ▪ searchFields - The list of fields that are relevant for the search by search term. Possible fields are NAME, DESCRIPTION, TAGS, AUTHOR. ▪ permissions - The list of required permissions. The EDIT value lists dashboards the user has an edit permission for. The EXECUTE value lists dashboards the user has an execute or view permission for. <p>Example</p> <pre>{ "keyword" : "mydashboard", "timestamp" : 1421135782252, "searchFields" : ["NAME", "DESCRIPTION", "TAGS", "AUTHOR"], "permissions" : ["EDIT", "EXECUTE"], "rating" : 3 }</pre>
MailService	Contains methods used to manage and send email notifications for sharing.
ManagementService	Contains methods to find license and statistics information for the Admin Console.
PolicyService	Contains methods to find and manage permission sets for run permissions, other artifact actions and access to specific features (entitlements) in MashZone NextGen Hub and the AppDepot. Also contains methods to create and manage dynamic groups for custom policies and manage policy configuration.
UserManagerService	Contains methods to find and manage users, groups and MashZone NextGen users attributes.
AliasService	Contains methods to manage MashZone NextGen alias definitions including create, delete, and assign alias definitions.
ACLService	Contains methods to manage permissions for dashboards, data feeds and aliases. It is possible store, delete, or assign permissions for mentioned definitions.

9.4.3 Input Parameters for JUMP Requests

JUMP requests contain a **params** property that takes either an array of parameters or an object with named parameters to pass to mashables, mashups or to platform service APIs. There are several ways to construct this array to accommodate different requirements:

- **No parameters** = an empty array, such as:
- **params: []**

- **Unnamed, ordered parameters with simple types** = array entries are separate by commas. They can be any legal JSON value or a mapping expression identifying a MashZone NextGen attribute, such as:

```
params: [ "short", 10 ]
```

or

```
params: [ "$global/AmazonId", true ]
```

- See Parameter Values (page 440) for details.
- **Named parameters with simple types in any order** = an object with matching properties for each named parameter, such as:

```
params: { "itemsPerPage": 10, "username": "$user/rssUserID" }
```

- The values for individual parameters (each property value) can be any legal JSON value or a mapping expression identifying a MashZone NextGen attribute. See Parameter Values (page 440) for details.
- Parameter names can also use namespaces, where required, as shown in the (following) example for a complex parameter.
- **A single, unnamed, complex parameter with named properties** = an object with any level of complexity and named properties.
- This is a special case of an unnamed parameter that is represented by a complex object with named properties. This example contains a complex object where properties may also have objects as values. Property names also include a namespace prefix:

```
params: { "ns:ItemLookup": { "ns:SubscriptionId": "0525E2PQ81DD7ZTWTk82",  
"ns:Validate": "False", "ns:Request": { "ns:ItemId": "B000EXS1BS",  
"ns:IdType": "ASIN", "ns:ResponseGroup": "Large" } } }
```

9.4.3.1 Parameter Values

Explicit parameter values, either named or unnamed, can be any type valid in JSON:

- **String** = characters enclosed in double quote marks, such as "string".
- **Number** = numeric characters with no quote marks, such as 104.2.
- **Keywords** = **true**, **false** or **null**
- **Arrays** = enclosed in brackets with commas separating array entries, such as [10, 5, 0].
- **Objects** = enclosed in braces with named properties. Property names are strings and end in a colon followed by the property value. For example, { propertyA: true } or { "ns:searchType": "book" }.

Note: Property names that include XML namespaces must be enclosed in double quote marks to ensure that the colon for the namespace is clear.

See the <http://www.json.org/json.org> site for a good introduction to JSON types and syntax.

Parameter values can also be resolved by the MashZone NextGen Server from a MashZone NextGen attribute. The value passed in the request is a mapping expression that identifies the MashZone NextGen attribute to use to resolve this input parameter.

9.5 Root Cause Miner

PPM provides the **Root Cause Miner** widget to analyze the visible data on a dashboard.

If you observe unusual symptoms on a dashboard, that is, interesting data points that need to be investigated, you can use the **Root Cause Miner** widget to analyze these symptoms. For example, you observe that the number of complaints in some distribution regions is too high and you want to get to the bottom of the symptom.

9.5.1 Symptoms and root causes

Symptoms are observed when a process analysis is performed, whereby the processes are analyzed with the help of a base set of filters.

Example

For example, you observe your global sales processes of the last year and you find that a disproportionately larger number of customer complaints occur in your subsidiary in **China** than in the rest of the world. You start a process analysis to find out the cause of this asymmetry. In addition to the filters that define the **base set** of your observation (here: sales processes, last year's processes, and country) you specify the dimension value that you want an explanation for (here: complaints). This dimension value is called the **symptom**. PPM uses your settings to analyze all suitable dimensions for values with strikingly high or low numbers of occurrences in combination with the symptom. The result of the analysis is a table that lists all dimension values together with an estimate of how strongly this value affects the occurrence of the symptomatic value. This table may contain a row showing that complaints very often occur together with the value "PN4711" for the product dimension. The dimension value "PN4711" is called a **root cause**.

The following criteria can be used as symptoms.

- Single-level dimension value, as mentioned in the example
- Multi-level dimensions, which are interpreted disjunctively, such as the dimension **Equipment** with the values **basic** or **luxury**.
- KPIs
- Value ranges, for example, **Processing time** with the range **10 to 12 hours**. When the symptom contains a process measure, all process instances for which that measure is undefined are counted among the non-symptomatic processes.

The following criteria are not supported.

- Time range dimensions are not supported.
- Function and relation dimensions. The symptom must be specified on process level.

SUPPORTED DIMENSIONS

- Only single-level text dimensions on process level and the variant dimension are analyzed as root causes.

Variants are analyzed on combined and precise level independently.

- Only dimensions that are available in the Process Mining context are included in the analysis.

NON-SUPPORTED DIMENSIONS

- Function and relation dimensions are not analyzed.
- Any dimension that already occurs in the symptom is ignored.
- Dimensions that contain user-defined steps or usages are ignored.

9.5.2 Limitations

Note the following limitations.

- Number of dimension values

PPM calculates only dimensions with less than a critical number of values in a query. Due to the risk of memory overflow, dimensions with more values are not calculated.

Normally there is no reason to analyze such dimensions. The typical use case is that of a search dimension, such as a dimension holding an order number, with one distinct value for each process instance.

- Number of dimension value occurrences

Dimension values that only occur in a few process instances are ignored in the calculation.

- Number of results

The maximal number of root causes that is returned to the client can be specified by the client as part of the query. If nothing is specified, PPM returns at most 50 results, consisting of the top 25 causes that promote the symptom and the top 25 causes that inhibit the symptom.

9.5.3 Score

The score output by PPM is a measure of the association between a process having the symptom and also having some other dimension value. This measure is a number between 1 and 1000. The value of the measure represents the importance of the root cause. The closer the value moves to 1000, the stronger the effect of the root cause becomes.

An association may promote the symptom, meaning that the root cause and the symptom extraordinarily often occur together, or may inhibit the symptom, meaning that they seldom occur together. In the example (page 441) above, a customer coming from China is associated with having a complaint, so the root cause **China** promotes the occurrence of complaints.

Note that the score does not necessarily correspond to the percentage of symptomatic processes. It takes other factors besides the percentage into account, such as the relative frequencies of the dimension values.

9.5.4 Scaled system

The **Root Cause Miner** feature is also available in a scaled system. The subservers transfer counts of their symptomatic and non-symptomatic processes for each dimension value to the master. The master aggregates these results, and computes the root causes and their scores. Note that this architecture may lead to high memory load on the subservers and much network traffic between the master and its subservers.

9.5.5 Permissions

No special permissions are required for performing root cause analysis. However, as usual your process and data access privileges in PPM restrict the data that you can analyze.

9.5.6 How to use the Root Cause Miner widget

For details on installing and using the **Root Cause Miner** widget, see the documentation **PPM widgets for MashZone NextGen**.

10 Additional Information and Support

10.1 Samples, Help and Other Documentation

You can find samples, help and other documentation for MashZone NextGen in:

- MashZone NextGen User and Developer Guide: You can download MashZone NextGen documentation at <https://empower.softwareag.com/> for offline access. An Empower account is required.
- **Online Help:** is accessible in MashZone NextGen from ? help buttons.
- **Documentation for Other Software AG Products:** is available online at <http://documentation.softwareag.com> with an Empower or from <http://techcommunity.softwareag.com/welcome-documentation> with a Software AG Tech Communities account.

10.2 Version and License Information

You can find version and other system information for MashZone NextGen in the Admin Console.

For system information:

1. Open the MashZone NextGen Admin Console (page 95).
2. Expand the **Server** menu section.
3. Click **System Info** and select the **Platform Information** or **System Information** tabs.
4. Click **License Info** for MashZone NextGen license information.

11 Legal information

11.1 Documentation scope

The information provided describes the settings and features as they were at the time of publishing. Since documentation and software are subject to different production cycles, the description of settings and features may differ from actual settings and features. Information about discrepancies is provided in the Release Notes that accompany the product. Please read the Release Notes and take the information into account when installing, setting up, and using the product.

If you want to install technical and/or business system functions without Software AG's consulting services, you require extensive knowledge of the system to be installed, its intended purpose, the target systems, and their various dependencies. Due to the number of platforms and interdependent hardware and software configurations, we can only describe specific installations. It is not possible to document all settings and dependencies.

When you combine various technologies, please observe the manufacturers' instructions, particularly announcements concerning releases on their Internet pages. We cannot guarantee proper functioning and installation of approved third-party systems and do not support them. Always follow the instructions provided in the installation manuals of the relevant manufacturers. If you experience difficulties, please contact the relevant manufacturer.

If you need help installing third-party systems, contact your local Software AG sales organization. Please note that this type of manufacturer-specific or customer-specific customization is not covered by the standard Software AG software maintenance agreement and can be performed only on special request and agreement.

11.2 Data protection

Software AG products provide functionality with respect to processing of personal data according to the EU General Data Protection Regulation (GDPR).

Where applicable, appropriate steps are documented in the respective administration documentation.