



USER AND DEVELOPER GUIDE

MASHZONE NEXTGEN

VERSION 10.7

April 2021

This document applies to MashZone NextGen Version 10.7 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 - 2021 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 <https://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 <https://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 <https://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 Open 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 8
 - 5.1.5 Pause automatic data refreshing 8
 - 5.1.6 Change column width and sort order 9
 - 5.1.7 Use multiple selection in lists and tables..... 9
 - 5.1.8 Save widget data as a CSV file..... 10
 - 5.1.9 Set the step width of a distribution chart..... 11
 - 5.1.10 Use bookmarks..... 11
 - 5.1.10.1 Open the bookmark bar 11
 - 5.1.10.2 Create bookmarks 12
 - 5.1.10.3 Open a bookmarked dashboard 12
 - 5.1.10.4 Rename a bookmark..... 12
 - 5.1.10.5 Delete a bookmark 13
 - 5.1.10.6 Update a bookmark 13
 - 5.1.10.7 Share a bookmark..... 14
 - 5.2 Use dashboard in edit mode 14
 - 5.2.1 Open a dashboard in the dashboard editor..... 14
 - 5.2.2 Create a dashboard in the dashboard editor 14
 - 5.2.3 Use smart dashboard edit mode..... 15
 - 5.2.3.1 Insert new row 15
 - 5.2.3.2 Delete a row..... 16
 - 5.2.3.3 Resize a row..... 16
 - 5.2.3.4 Lock row height 16
 - 5.2.3.5 Change a row order 16
 - 5.2.3.6 Place a widget 16
 - 5.2.3.7 Resize widgets..... 17
 - 5.2.3.8 Avoid line break in dashboard view mode 17
 - 5.2.3.9 Group widgets 17
 - 5.2.3.10 Set row style 18
 - 5.2.3.11 Set layout group style..... 18
 - 5.2.3.12 Set dashboard tab style 18
 - 5.2.3.13 Switch to fixed-grid workspace..... 18
 - 5.2.4 Save a dashboard 19
 - 5.2.5 Edit a dashboard..... 19
 - 5.2.6 Delete a dashboard 20
 - 5.2.7 Edit dashboard properties 20
 - 5.2.8 Manage dashboard permissions..... 20
 - 5.2.9 Change the dashboard style 21
 - 5.2.10 Process Mining context 22
 - 5.2.10.1 Create a Process Mining context 22
 - 5.2.10.2 Edit a Process Mining context..... 24
 - 5.2.10.3 Delete a Process Mining context 25

5.2.10.4	Assign data columns to Process Mining context-based widgets	25
5.2.10.5	Configure filters for Process Mining context-based widgets	26
5.2.10.5.1	Configure a date filter for a Process Mining context	27
5.2.10.6	Exclude a widget from a Process Mining context filter	28
5.2.10.7	Enable the filter panel	29
5.2.10.8	Configure the filter panel	29
5.2.10.9	Migrate a dashboard context	30
5.2.11	Use Multi-lingual dashboards	31
5.2.11.1	Configure languages for translation	31
5.2.11.2	Translate the widget content	32
5.2.11.3	Select the language for your dashboard	32
5.2.11.4	Translate the tab title	33
5.2.12	Manage bookmarks	33
5.2.12.1	Enable the bookmark bar	33
5.2.12.2	Open the bookmark bar	34
5.2.12.3	Rename a bookmark	34
5.2.12.4	Delete a bookmark	34
5.2.12.5	Share a bookmark	35
5.2.12.6	Set and reset a bookmark as default	35
5.3	Configure widgets	36
5.3.1	Set up widgets	37
5.3.2	Insert widgets in a dashboard	37
5.3.3	Copy and paste widgets in dashboards	37
5.3.4	Assign data sources to widgets	38
5.3.4.1	Assign data columns to line, column or bar charts	39
5.3.4.1.1	Set properties of widget elements	39
5.3.4.1.2	Configure KPI thresholds	42
5.3.4.1.3	Configure KPI rating	42
5.3.4.2	Assign data columns to a distribution chart	43
5.3.4.2.1	Set properties of widget elements	44
5.3.4.3	Assign data columns to pie charts	46
5.3.4.3.1	Set properties of widget elements	46
5.3.4.4	Assign data columns to bubble charts	47
5.3.4.4.1	Set properties of widget elements	48
5.3.4.4.2	Configure KPI thresholds	49
5.3.4.5	Assign data columns to grids	50
5.3.4.5.1	Set properties of widget elements	50
5.3.4.5.2	Configure thresholds	52
5.3.4.6	Assign data columns to lists	53
5.3.4.6.1	Set properties of widget elements	54
5.3.4.7	Assign data columns to circular gauge charts	55
5.3.4.7.1	Set properties of widget elements	55
5.3.4.7.2	Configure thresholds	56
5.3.4.8	Assign data columns to horizontal and vertical gauge charts	57
5.3.4.8.1	Set properties of widget elements	57
5.3.4.8.2	Configure thresholds	58
5.3.4.9	Assign data columns to traffic lights	59
5.3.4.9.1	Set properties of widget elements	59
5.3.4.9.2	Configure thresholds	60
5.3.4.10	Assign data columns to a drop-down box	61

5.3.4.10.1	Set properties of widget elements.....	61
5.3.4.11	Assign data columns to labels	62
5.3.4.11.1	Set properties of widget elements.....	62
5.3.4.12	Assign data columns to a rich text area	63
5.3.4.12.1	Set properties of widget elements.....	64
5.3.4.13	Assign data columns to sliders	64
5.3.4.13.1	Set properties of widget elements.....	64
5.3.4.14	Assign data columns to date filters	65
5.3.4.15	Assign data columns to vector map.....	65
5.3.4.15.1	Set properties of widget elements.....	66
5.3.4.15.2	Configure KPI thresholds	67
5.3.4.16	Assign data columns to map with markers	68
5.3.4.16.1	Set properties of widget elements.....	68
5.3.4.16.2	Configure KPI thresholds	69
5.3.4.17	Create input parameters	70
5.3.4.18	Calculate the feed data	71
5.3.4.19	Specify the color mapping for dimensions	72
5.3.5	Display legend.....	73
5.3.6	Remove widgets from dashboard	74
5.3.7	Define filters for widgets	74
5.3.8	Specify actions for widgets	76
5.3.8.1	Change tab	76
5.3.8.2	Set selection.....	77
5.3.8.3	Call URL.....	77
5.3.8.4	Specify an image source URL	78
5.3.8.5	Post data	79
5.3.8.6	Send events	80
5.3.9	Hide or display widget header and border.....	81
5.3.10	Resize widgets.....	82
5.3.11	Move widgets to the foreground or background	82
5.3.12	Display widgets on multiple dashboard tabs.....	82
5.3.13	Change the widget style.....	83
5.4	Configure dashboard tabs	83
5.4.1	Add dashboard tabs	84
5.4.2	Delete dashboard tabs.....	84
5.4.3	Set dashboard tab properties.....	84
5.4.4	Set dashboard tab style	85
5.4.5	Set dashboard tab as default	85
5.5	Other.....	85
5.5.1	Use dynamic URL selection	85
5.5.2	Hide tab bar in view mode.....	87
5.5.3	Set display size	87
5.5.4	Set grid lines.....	87
5.5.5	Allow anonymous access to dashboards	88
5.5.5.1	PPM data source operator.....	88
5.5.5.2	File resources	88
5.5.6	Display dashboards without application header.....	89
5.5.7	Export and import dashboards	89
5.5.8	Customize styles and themes	89

6	Use the data feed editor	91
6.1	Create a data feed	91
6.1.1	Open the data feed editor	91
6.1.2	Select data sources	91
6.1.3	Calculate the feed data	92
6.1.4	Add further operators	93
6.1.5	Connect the operators	94
6.1.6	Save the data feed	94
6.2	Edit a data feed	95
6.3	Delete data feeds	95
6.4	Edit data feed properties	95
6.5	Manage data feed permissions	96
6.6	Data feed operators	97
6.7	Use Materialized Feeds	97
6.7.1	Use datasets storing data feed results	98
6.7.2	Manage Materialized Feeds in Terracotta	98
6.7.3	Use datasets from an external application	99
6.7.4	Which privileges are required?	100
6.8	Other	100
6.8.1	Export and import data feeds	100
6.8.2	Display data feed without application header	100
7	Good to know	103
7.1	Dashboards	103
7.2	Data feeds	103
7.3	What is a view-only license?	104
8	Administration	106
8.1	Open the Admin Console	106
8.2	MashZone NextGen security	106
8.2.1	Manage your MashZone NextGen profile	107
8.2.1.1	Manage your locale and account information	107
8.2.1.2	Change your password	107
8.2.2	MashZone NextGen server	108
8.2.3	MashZone NextGen repository	108
8.2.4	Authentication and Guest Access	108
8.2.4.1	User Authentication	109
8.2.4.2	Valid Credentials	109
8.2.4.3	Sessions and Timeouts	110
8.2.5	Default user account	110
8.2.6	Authentication with single sign-on solutions	110
8.2.6.1	Configuration for Agent-Based SSO Solutions	111
8.2.6.2	Implementing a Custom SSO Filter	112
8.2.6.3	SSO integration in My webMethods	113
8.2.7	Protect MashZone NextGen Event Service access	114
8.2.8	Anti-Clickjacking prevention when using iFrame	114
8.2.8.1	MashZone NextGen HTTP header security filter	115
8.2.8.1.1	Example	116
8.2.8.2	MashZone NextGen Content Security Policy	116
8.2.8.3	Add a trusted site to allow iFrame	116
8.2.8.4	Add multiple trusted sites to allow iFrame	117
8.2.8.5	Content-Security-Policy using wildcards	117
8.2.9	Handle personal data in log files	117

8.2.10	Create whitelists for URL calls.....	118
8.3	Getting Started with the MashZone NextGen Server	119
8.3.1	Additional MashZone NextGen System and Software Requirements	120
8.3.1.1	Additional Recommendations for MashZone NextGen.....	120
8.3.2	What is Installed with MashZone NextGen	121
8.3.3	Start and Stop the MashZone NextGen Server	121
8.3.3.1	Start the MashZone NextGen Event Service	121
8.3.3.2	Start the MashZone NextGen Server	122
8.3.3.3	Stop the MashZone NextGen Event Service.....	122
8.3.3.4	Stop the MashZone NextGen Server.....	123
8.3.4	Startup Considerations	123
8.3.5	Manage Licenses for MashZone NextGen and BigMemory Max	123
8.3.6	Move the MashZone NextGen repository to a robust database solution	124
8.3.6.1	Troubleshooting Connections to the MashZone NextGen Repository.....	125
8.3.6.2	Move the MashZone NextGen repository to Microsoft SQL Server ...	126
8.3.6.3	Move the MashZone NextGen repository to MySQL.....	128
8.3.6.4	Move the MashZone NextGen repository to Oracle.....	129
8.3.6.5	Move the MashZone NextGen repository to PostGres	131
8.3.7	Integrate Your LDAP Directory with MashZone NextGen.....	132
8.3.7.1	Defining LDAP Connection Configuration	133
8.3.7.2	Defining the Authentication Scheme.....	134
8.3.7.3	Defining the Authorization Scheme	135
8.3.7.4	Enabling MashZone NextGen Application Queries for All LDAP Users or Groups for Permissions	136
8.3.7.5	Start MashZone NextGen with an initial administrator user.....	138
8.3.8	Use the default MashZone NextGen user repository	138
8.3.8.1	Manage users	139
8.3.8.1.1	Create users	139
8.3.8.1.2	Edit, assign roles and other user management tasks.....	139
8.3.8.2	Manage user groups	140
8.3.8.3	Which user roles exist?	141
8.3.9	Install MashZone NextGen and MashZone NextGen Event Service as Windows services.....	141
8.3.10	Command Central	141
8.3.10.1	MashZone NextGen plug-in.....	142
8.3.10.1.1	Instance Overview	142
8.3.10.1.2	Instance Configuration	142
8.3.10.1.3	Instance Logs	144
8.3.10.1.4	Enable remote JMX connection	144
8.3.10.2	MashZone NextGen Event Service plug-in	145
8.3.10.2.1	Instance overview	145
8.3.10.2.2	Instance configuration.....	145
8.3.10.2.3	Instance Logs	147
8.3.10.3	Configure ports.....	147
8.3.10.4	Configure keystores.....	148
8.3.10.5	Configure truststores	149
8.3.10.6	Configure TLS certificates	149
8.3.10.7	Manage licenses.....	149
8.3.11	Required installer modules	150
8.4	MashZone NextGen Server Configuration	150
8.4.1	Memory Configuration for the MashZone NextGen Server	150

8.4.1.1	Configuration When MashZone NextGen Uses Only Heap Memory...	151
8.4.1.2	Configuration When MashZone NextGen Uses Heap and Off-Heap Memory	152
8.4.2	Configure feed processing time zone.....	153
8.4.3	Support International Character Sets and Locales.....	154
8.4.3.1	Set the Repository Character Encoding	154
8.4.3.2	Set the Repository Timezone or Offset.....	155
8.4.3.3	Date, Time and Numeric Display Options	155
8.4.3.4	Message Log and Default Locales.....	155
8.4.4	Configure the MashZone NextGen server with custom ports	156
8.4.4.1	Change MashZone NextGen Server Ports	156
8.4.4.2	Change MashZone NextGen Repository Ports.....	156
8.4.4.3	Tomcat Application Server Port	157
8.4.5	Configure the MashZone NextGen server to work with a proxy server.....	157
8.4.6	Embedding MashZone NextGen in external system environments.....	157
8.4.6.1	Configure MashZone NextGen server to work with iFrame	158
8.4.6.2	Post data	159
8.4.6.3	URL selection.....	159
8.4.7	Define a Proxy Server White list for MashZone NextGen	159
8.4.7.1	Using Regular Expressions in a White list	160
8.4.7.2	Specifying Literal Dot Separators.....	160
8.4.7.3	Specifying Domains	161
8.4.7.4	Specifying Host Names	161
8.4.8	Configure MashZone NextGen for TLS and Digital Certificates.....	161
8.4.8.1	The Certificate Store and Certificates	162
8.4.8.2	Key Certificate Pairs	162
8.4.8.3	Trusted Peer Certificates	162
8.4.8.4	The Certificate Store.....	163
8.4.8.5	Configure Mutual TLS Between Users and MashZone NextGen.....	163
8.4.8.6	One-Way TLS to MashZone NextGen	163
8.4.8.7	One-Way TLS to Information Sources.....	163
8.4.8.8	Configure HTTPS and Certificate Stores in the Application Server ...	164
8.4.8.9	Update TLS Configuration for Java	165
8.4.9	MashZone NextGen Logging	166
8.4.9.1	Configure Logging for the MashZone NextGen Server	166
8.4.9.2	Audit logging for dashboards, data feeds, aliases, and permissions.....	167
8.4.10	BigMemory Max for Caching, Connections and In-Memory Stores.....	168
8.4.10.1	Caching for the MashZone NextGen Server	169
8.4.10.1.1	Artifact Caching.....	169
8.4.10.1.2	Response Caching	169
8.4.10.1.3	Distributed Caching for MashZone NextGen Clusters.....	169
8.4.10.1.4	Configure BigMemory Max Servers for MashZone NextGen Caching	170
8.4.10.2	Working with BigMemory Max Stores for RAQL	171
8.4.10.2.1	Declared In-Memory Stores	172
8.4.10.2.1.1	Declare a new In-Memory Store.....	172
8.4.10.2.1.2	Modify a Declared In-Memory Store.....	174
8.4.10.2.1.3	View Details for Declared In-Memory Stores	174
8.4.10.2.1.4	Define permissions for declared In-Memory Stores	175
8.4.10.3	Dynamic In-Memory Stores	175

8.4.10.3.1	Manage Dynamic BigMemory Max Stores for MashZone NextGen Analytics	176
8.4.10.3.2	Add an External Dynamic In-Memory Store Connection	176
8.4.10.3.3	Delete External Dynamic In-Memory Store Connections	177
8.4.10.3.4	Define permissions for external dynamic In-Memory Store Connections	178
8.4.11	Manage Terracotta connections	178
8.4.11.1	Register Terracotta connections	178
8.4.11.2	Edit Terracotta connections	179
8.4.11.3	Test Terracotta connections	180
8.4.11.4	Share Terracotta connections	180
8.4.11.5	Delete Terracotta connections	181
8.4.11.6	Ensure secure interaction with Terracotta	181
8.4.12	Manage data sources and drivers	182
8.4.12.1	Add a data source	182
8.4.12.2	Edit, test or remove data sources	184
8.4.12.3	Share data sources	184
8.4.12.4	Add or manage JDBC drivers	185
8.4.12.5	Migrate JDBC connections	185
8.4.12.5.1	Migrate JDBC configuration of Presto to MashZone NextGen ..	186
8.4.12.5.2	Migrate JDBC connections of Presto to MashZone NextGen	186
8.4.12.5.3	Migrate JDBC configuration of MashZone NextGen 9.10	187
8.4.12.5.4	Migrate JDBC connections of MashZone legacy to MashZone NextGen	188
8.4.13	Manage geographical map resources	188
8.4.13.1	Manage geoJSON files	188
8.4.13.2	Manage tile server configuration files	189
8.4.14	Tune memory/caching for MashZone NextGen	189
8.4.14.1	Tune MashZone Memory and Cache Configuration Manually	189
8.4.14.2	Update Cache Memory Settings	190
8.4.14.3	Update MashZone ThreadSize Properties	190
8.4.15	Manage Materialized Feeds	190
8.4.15.1	Configure Materialized Feeds	190
8.4.15.2	Edit Materialized Feeds	191
8.4.15.3	Delete Materialized Feeds	191
8.4.15.4	Share Materialized Feeds	192
8.4.15.5	Materialize data feed results in Terracotta	193
8.4.15.5.1	Use the Admin Console	193
8.4.15.5.2	Use the server administration tool	193
8.4.15.5.3	Use server web services	194
8.4.16	Disable automatic masking in CSV files	195
8.5	MashZone NextGen Server Administration	196
8.5.1	Manage files for MashZone NextGen dashboards and data feeds	196
8.5.1.1	Add external resources as MashZone NextGen files	196
8.5.1.2	Find MashZone NextGen files	197
8.5.1.3	Update or delete MashZone NextGen files	197
8.5.1.4	Share MashZone NextGen resource files	198
8.5.2	Manage resource directories	198
8.5.2.1	Create resource directory	198
8.5.2.2	Change resource directory	199
8.5.2.3	Delete resource directory	199
8.5.2.4	Share resource directory	199

8.5.3	Manage URL aliases.....	200
8.5.3.1	Create URL alias.....	200
8.5.3.2	Change URL alias	201
8.5.3.3	Delete URL alias.....	201
8.5.3.4	Share URL alias.....	201
8.5.4	Deploying MashZone NextGen instances, clusters, or artifacts	202
8.5.4.1	Deploying the core widgets.....	202
8.5.4.2	Deploying artifacts and other metadata using the command line	202
8.5.4.2.1	Export users, groups, and role assignments	204
8.5.4.2.2	Export role assignments for users and groups.....	205
8.5.4.2.3	Export dashboards.....	205
8.5.4.2.4	Export data feeds	206
8.5.4.2.5	Export aliases	207
8.5.4.2.6	Import users, groups, and role assignments.....	208
8.5.4.2.7	Import role assignments for users and groups.....	209
8.5.4.2.8	Import dashboards	210
8.5.4.2.9	Import data feeds.....	211
8.5.4.2.10	Import aliases.....	211
8.5.4.2.11	Deploying multiple MashZone NextGen servers in one host....	212
8.5.4.3	Deploying artifacts using the Admin Console	212
8.5.4.3.1	Export dashboards, data feeds, and aliases	212
8.5.4.3.2	Import dashboards, data feeds, and aliases.....	213
8.5.5	Clustering MashZone NextGen Servers.....	214
8.5.5.1	Setting Up a New Cluster	214
8.5.5.2	Adding New Members to an Existing Cluster.....	216
8.5.6	Sharing the MashZone NextGen Repository in Clustered Environments	217
8.5.6.1	Create and Share a New MashZone NextGen Repository	217
8.5.6.2	Share an Existing MashZone NextGen Repository.....	217
8.5.7	Setting Up an External MashZone NextGen Configuration Folder	218
8.5.7.1	MashZone NextGen File-Based Configuration and Extensions.....	219
8.5.7.2	MashZone NextGen Configuration Files That Can Be External.....	219
8.5.7.3	MashZone NextGen Configuration Files That Must Be Internal	220
8.5.7.4	MashZone NextGen Extensions.....	221
8.5.8	MashZone NextGen dashboards in a clustered scenario	222
8.5.8.1	Preliminary.....	222
8.5.8.2	Configuration.....	222
8.5.8.2.1	MashZone NextGen Event Service.....	222
8.5.8.3	rtbs.properties.....	223
8.5.8.4	zookeeper.properties	223
8.5.8.5	server.properties.....	224
8.5.8.6	MashZone NextGen nodes	224
8.5.8.6.1	Customizing dashboards.....	225
8.5.8.7	Custom styles	225
8.5.8.8	Custom widgets	225
8.5.8.9	Using JDBC drivers.....	225
8.5.8.10	Local file resources.....	226
8.5.9	Customize application and dashboard styles	226
8.5.9.1	Customize dashboard and widget styles.....	226
8.5.9.2	Customize the application style	227
8.5.9.3	Customize the MashZone NextGen welcome page	227
8.5.9.4	Download styles.....	228
8.5.9.5	Upload styles.....	229

8.5.9.6	Compile styles	230
8.5.9.7	Delete styles	230
8.5.9.8	Switch the style editing mode	231
8.5.9.9	Adjust the font size to different devices	231
8.5.9.10	Custom widgets	233
8.5.9.11	How are the style files structured?.....	233
8.5.10	Empty MashZone NextGen caches	235
8.6	Event Service Configuration and Administration	236
8.6.1	About the Event Service and Event Data	236
8.6.2	Use Events as Information Sources.....	236
8.6.3	Manage EDA Event Sources.....	237
8.6.3.1	Create EDA Event Sources.....	237
8.6.3.2	Edit EDA Event Sources.....	241
8.6.3.2.1	Simple and Hybrid Throttling Strategies.....	244
8.6.3.3	Duplicate EDA Event Sources	245
8.6.3.4	Delete EDA Event Sources	246
8.6.3.5	Share EDA Event Sources.....	246
8.6.4	Manage Apama Event Sources	247
8.6.4.1	Create Apama Event Sources	247
8.6.4.2	Duplicate Apama Event Sources	252
8.6.4.3	Delete Apama Event Sources	252
8.6.4.4	Edit Apama Event Sources.....	253
8.6.4.5	Share Apama Event Sources.....	258
8.6.5	Start or Stop an Event Source	258
8.6.6	Restart all Event Sources	259
8.6.7	Manage Apama Instances	259
8.6.7.1	Create Apama Instances	260
8.6.7.2	Edit Apama Instances	260
8.6.7.3	Delete Apama Instances.....	261
8.6.8	Manage Apama Event Targets.....	261
8.6.8.1	Create Apama Event Targets.....	261
8.6.8.2	Edit Apama Event Targets	262
8.6.8.3	Delete Apama Event Targets.....	263
8.6.8.4	Share Apama Event Target	263
8.7	Process Performance Manager Integration	264
8.7.1	Manage PPM Connections	264
8.7.2	Create PPM Connections.....	264
8.7.3	Edit PPM Connections	265
8.7.4	Delete PPM Connections.....	266
8.7.5	Share PPM connections	267
8.8	webMethods Business Console Integration.....	267
8.8.1	Example	268
8.8.2	Authentication	268
8.8.3	Example URL.....	268
8.8.4	Configuration	269
8.8.5	Outbound API.....	270
8.8.6	Inbound API.....	270
8.9	MashZone NextGen Repositories	270
8.9.1	Maintenance Suggestions.....	271
8.9.2	Tuning the MashZone NextGen Repository Connection Pool	271
8.9.2.1	Connection Pool Size Properties.....	271
8.9.2.2	Idle Pool Connection Properties	271

8.9.3	Synchronize the MashZone NextGen Repository and MashZone NextGen Server Time Zones.....	272
9	Appendix.....	273
9.1	Widgets.....	273
9.1.1	Action button	273
9.1.2	Bar chart	273
9.1.3	Bubble chart	275
9.1.4	Circular gauge chart.....	276
9.1.5	Column chart	277
9.1.6	Date filter	279
9.1.7	Distribution chart.....	281
9.1.8	Drop-down box.....	282
9.1.9	Grid	283
9.1.10	Horizontal and vertical gauge chart	285
9.1.11	Image	286
9.1.12	Input field.....	287
9.1.13	Label.....	288
9.1.14	Line chart	289
9.1.15	List.....	290
9.1.16	Map with markers	292
9.1.17	Pie chart.....	293
9.1.18	Rich text area.....	294
9.1.19	Slider	296
9.1.20	Jump to PPM client.....	299
9.1.21	Traffic lights.....	300
9.1.22	Vector map	301
9.2	Operators.....	302
9.2.1	Data source operators	303
9.2.1.1	ARIS table.....	303
9.2.1.2	BigMemory Max	304
9.2.1.3	CSV	306
9.2.1.4	Data feed.....	308
9.2.1.5	Events	308
9.2.1.6	JDBC	309
9.2.1.7	JSON.....	310
9.2.1.8	Excel	312
9.2.1.9	PPM.....	315
9.2.1.9.1	Add filters	316
9.2.1.10	Terracotta	317
9.2.1.11	XML	320
9.2.1.12	URL parameter syntax	321
9.2.1.13	RAQL and SQL statement parameters.....	322
9.2.2	Data transformation operators	323
9.2.2.1	Aggregate	323
9.2.2.2	Arithmetic	326
9.2.2.3	Average	327
9.2.2.4	Change data type.....	328
9.2.2.5	Change data type - single value	333
9.2.2.6	Column to value.....	339
9.2.2.7	Combine data feeds.....	339
9.2.2.8	Concatenate data feeds.....	340
9.2.2.9	Concatenate texts	340
9.2.2.10	Conditional replace.....	341
9.2.2.11	Convert text.....	343

9.2.2.12	Copy data feeds	344
9.2.2.13	Copy single value	344
9.2.2.14	Delete column	344
9.2.2.15	Duplicate column	345
9.2.2.16	Extract text	345
9.2.2.17	Filter by date	346
9.2.2.18	Filter rows	347
9.2.2.19	Find text index.....	349
9.2.2.20	Goal accomplishment.....	350
9.2.2.21	Insert column	351
9.2.2.22	Merge single texts	351
9.2.2.23	Move date	352
9.2.2.24	Named copy	353
9.2.2.25	RAQL Inline	354
9.2.2.26	Rename column	354
9.2.2.27	Replace text	355
9.2.2.28	Round up/down	356
9.2.2.29	Round up/down date	357
9.2.2.30	Runtime info.....	359
9.2.2.31	Value to column	360
9.2.3	User input operators.....	360
9.2.3.1	Date user input.....	360
9.2.3.2	Date user input (List).....	361
9.2.3.3	Number user input	362
9.2.3.4	Number user input (List)	363
9.2.3.5	Text user input	363
9.2.3.6	Text user input (List)	364
9.3	Analytics with RAQL	365
9.3.1	About the Real-Time Analytical Query Language.....	365
9.3.1.1	Scope of RAQL.....	365
9.3.1.2	RAQL in MashZone NextGen	365
9.3.1.3	Query Syntax of RAQL	366
9.3.1.4	Data types in RAQL	367
9.3.2	Language reference.....	368
9.3.2.1	Reserved keywords	368
9.3.2.2	Syntax Reference	370
9.3.2.2.1	ExplainOrSelectStatement	370
9.3.2.2.2	CharacterConstant.....	370
9.3.2.2.3	SelectStatement.....	370
9.3.2.2.4	SelectWithOrder	371
9.3.2.2.5	QueryExpressionBody	371
9.3.2.2.6	QueryTerm	371
9.3.2.2.7	QueryPrimary	372
9.3.2.2.8	SelectWithoutOrder	372
9.3.2.2.9	SelectList	372
9.3.2.2.10	SelectItem	373
9.3.2.2.11	SelectStar	373
9.3.2.2.12	ObjectName	373
9.3.2.2.13	SqlSimpleExpression	374
9.3.2.2.14	SqlMultiplicativeExpression	374
9.3.2.2.15	SqlExponentExpression	375
9.3.2.2.16	SqlUnaryExpression	375
9.3.2.2.17	SqlPrimaryExpression.....	376
9.3.2.2.18	SqlExpression	376
9.3.2.2.19	SqlAndExpression	377

9.3.2.2.20	SqlUnaryLogicalExpression.....	377
9.3.2.2.21	SqlRelationalExpression.....	378
9.3.2.2.22	SqlExpressionList	378
9.3.2.2.23	SqlRelationalOperatorExpression	378
9.3.2.2.24	Relop.....	379
9.3.2.2.25	SqlInClause	379
9.3.2.2.26	SqlBetweenClause	380
9.3.2.2.27	SqlLikeClause	380
9.3.2.2.28	IsNullClause	380
9.3.2.2.29	Constant	381
9.3.2.2.30	FunctionCall.....	382
9.3.2.2.31	CastableDataTypeDeclaration	384
9.3.2.2.32	UserDefinedType	385
9.3.2.2.33	SimpleCaseCall	385
9.3.2.2.34	BooleanCaseCall.....	385
9.3.2.2.35	AggregateFunctionCall.....	386
9.3.2.2.36	SourceField	387
9.3.2.2.37	GenericFunctionCall	387
9.3.2.2.38	ConstantList	387
9.3.2.2.39	WindowSpecification	388
9.3.2.2.40	OrderByClause	388
9.3.2.2.41	WindowFrameClause.....	388
9.3.2.2.42	WindowFrameStart	389
9.3.2.2.43	FrameBound	389
9.3.2.2.44	WindowFrameBound	389
9.3.2.2.45	GroupingOperation	390
9.3.2.2.46	WindowFunctionCall.....	390
9.3.2.2.47	LongConstant	391
9.3.2.2.48	IntegerConstant.....	391
9.3.2.2.49	ArrayElementReference	391
9.3.2.2.50	ArrayValueConstructor	391
9.3.2.2.51	Alias	392
9.3.2.2.52	FromClause	392
9.3.2.2.53	FromItem.....	392
9.3.2.2.54	SourceReferenceOrSubquery	392
9.3.2.2.55	SourceReference	393
9.3.2.2.56	JoinedTable	393
9.3.2.2.57	SourceFieldList.....	394
9.3.2.2.58	WhereClause	394
9.3.2.2.59	GroupByClause	394
9.3.2.2.60	GroupingElementList.....	394
9.3.2.2.61	GroupingElement.....	395
9.3.2.2.62	EmptyGroupingSet	395
9.3.2.2.63	OrdinaryGroupingSet	395
9.3.2.2.64	GroupingExpressionReferenceList.....	396
9.3.2.2.65	GroupingExpressionReference	396
9.3.2.2.66	RollupList	396
9.3.2.2.67	OrdinaryGroupingSetList.....	396
9.3.2.2.68	CubeList.....	397
9.3.2.2.69	GroupingSetsSpecification	397
9.3.2.2.70	GroupingSetList	397
9.3.2.2.71	GroupingSet	398
9.3.2.2.72	HavingClause.....	398
9.3.2.2.73	LimitClause.....	398
9.3.2.2.74	S_NUMBER.....	399
9.3.2.2.75	FLOAT.....	399
9.3.2.2.76	INTEGER	399

9.3.2.2.77	DIGIT	400
9.3.2.2.78	S_PARAMETER	400
9.3.2.2.79	S_IDENTIFIER.....	400
9.3.2.2.80	NAMESTARTCHAR.....	401
9.3.2.2.81	NAMECHAR.....	402
9.3.2.2.82	S_QUOTED_IDENTIFIER.....	402
9.3.2.2.83	S_CHAR_LITERAL	403
9.3.2.2.84	IgnorableWhitespace.....	403
9.3.2.2.85	LINE_COMMENT	403
9.3.2.2.86	MULTI_LINE_COMMENT.....	404
9.3.3	RAQL extensions.....	404
9.3.3.1	RAQL function extensions	404
9.3.3.1.1	Plain functions	404
9.3.3.1.2	Analytic functions	410
9.3.3.1.3	Window functions over RANGES	414
9.3.3.1.4	Migrate RAQL queries from Version 3.7 to 3.8.....	417
9.3.3.1.4.1	Update plain functions for 3.8.....	418
9.3.3.1.4.2	Update analytic functions for 3.8.....	419
9.3.3.2	Creating new user-defined functions	420
9.3.3.2.1	Set Up Your Development Environment	421
9.3.3.2.2	Write Plain Functions for RAQL	422
9.3.3.2.3	Configure, Compile, Deploy and Test User-Defined Functions	424
9.3.3.2.3.1	UDF Library Configuration	425
9.3.3.2.3.2	External UDF Library Deployment Folder	426
9.3.3.2.4	Write Window Analytic Functions for RAQL	427
9.3.3.2.4.1	Construct and Initialize the Window Analytic Function Class	427
9.3.3.2.4.2	Implement the call Method for the Window Analytic Function	428
9.3.3.2.4.3	Complete WindowSum Example.....	430
9.3.3.2.4.4	Work with Specific Records in Window Calculations.....	432
9.3.3.2.5	Write Aggregate Analytic Functions for RAQL.....	433
9.3.3.2.5.1	My Average Aggregate Example.....	434
9.3.3.2.5.2	Kurtosis Using a Third Party Library.....	436
9.3.3.2.6	Use a Factory for Function Overloading.....	439
9.4	Root Cause Miner	439
9.4.1	Symptoms and root causes	439
9.4.2	Limitations.....	440
9.4.3	Score	440
9.4.4	Scaled system	441
9.4.5	Permissions.....	441
9.4.6	How to use the Root Cause Miner widget.....	441
10	Additional Information and Support.....	442
10.1	Samples, Help and Other Documentation	442
10.2	Version and License Information	442
11	Legal information	443
11.1	Documentation scope.....	443
11.2	Support.....	443

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 of MashZone NextGen version 10.7.

TOPICS

Introducing Software AG MashZone NextGen

MashZone NextGen welcome page (page 3)

Open a dashboard in MashZone NextGen (page 4)

Use the dashboard editor (page 5)

Use the data feed editor (page 91)

Administration (page 106)

Appendix (page 273)

Additional Information and Support (page 442)

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 103) are represented graphically and analyzed on dashboards (page 103).

IN-MEMORY ANALYTICS

- MashZone NextGen leverages BigMemory Max 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 Max, 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 available in MashZone NextGen. Alternatively, you can use the **Dashboards** menu in the program bar to open any available dashboard.

To open an available data feed, you can use the **Data Feeds** menu in the program bar.

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.

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.

Prerequisites

You have the appropriate privileges to open a dashboard in view or edit mode (page 20).

You have the appropriate privileges to view the data of the data feeds (page 96) that are used by the dashboard.

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

You can use the graphical user interface of the dashboard editor to easily create, manage, and view 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 open the dashboard editor, the **Dashboard Designer** role is required. The **Dashboard Designer** role can be assigned to users in the user administration (page 138) group.

To display a dashboard in view mode, the **View** permission is required. To edit a dashboard in dashboard editor, the edit permission is required. For details, see Manage dashboard permissions (page 20).

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 20) for details.

5.1.1 Open 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.

Prerequisite

You have the required access privileges for the dashboard (page 20).

Procedure

1. Create (page 14) or open (page 14) 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.

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 277), Bar chart

(page 273), Pie chart (page 293), Grid (page 283), and List (page 290). The widget to be filtered must also support multiple selection. For details, see the chapter Define filters for widgets (page 74).

DISTRIBUTION CHARTS

In distribution charts, you can select multiple values using a selection frame by default. If you select a range of values using a selection frame, the upper end of the range is not included in the selection for filtering. For example, if you select a range from 10 to 20 and the step width is 1, the affected widget is filtered with the values 10 to 19.

Procedure

1. Open a dashboard in dashboard view mode. (page 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. In distribution charts, you can use a selection frame to select multiple columns.
 - a. Click a column and drag the mouse to the left or right.
 - b. A selection frame is displayed.
 - c. Release the mouse pointer.

The selected columns are highlighted.

4. 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 the chapter Define filters for widgets (page 74).

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 (page 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 a list.
- Numeric filter criteria are provided as a 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 296)** widget for details.

Prerequisites

The filter panel is enabled (page 29).

Procedure

1. Open a dashboard in dashboard view mode. (page 5)
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 25).
 - 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 26) 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.
6. Click the name of a text-based filter criterion to sort the values of the criterion.
7. Select the individual values of the text-based 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 (page 38) dialog.

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

Procedure

1. Open a dashboard in dashboard view mode. (page 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 reimplemented 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 (page 38) dialog.

Procedure

1. Display a dashboard in dashboard view mode. (page 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. (page 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 multiple selection in lists and tables

You can use the multiple selection of values to filter values of other widgets (page 5).

In lists and tables, you can select individual or multiple values using your mouse and keyboard. A list also provides check boxes to select multiple values if the **Multiple selection** option is enabled. For details, see the option list of the List (page 290) and Grid (page 283) widgets.

Procedure

1. Open a dashboard in view mode. (page 5)
2. Click a row in a list or table to select the value contained in it.
3. To select multiple values, click a row, press the **Shift** key and click another row.

All other rows between the selected rows are now also selected.
4. Alternatively, hold down the **Ctrl** key and click individual rows in the list or table to select multiple values.
5. To select all values in a list, click the check box in the column header. Click the check box again to clear all list values.
6. To sort a list or table, click the name in the header of a column.

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.

7. 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 277), Bar chart (page 273), and Pie chart (page 293).

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 aggregation and the value rounding are taken from the column configuration set in the data assignment (page 38) (Assign data 2/2 dialog). The sorting is taken from the data feed result based on the feed definition.

Procedure

1. Open a dashboard in dashboard view mode. (page 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 if required.
 - a. Select a separator character from the **Separator** drop-down menu. The default separator is ,.
 - b. Select a masking character from the **Masking** drop-down menu. The default masking is ".
 - c. Select an encoding type from the **Encoding** drop-down menu. The default encoding is **UTF-8**.
 - d. Enable the **Define formatting** option to change the default data format. You can select a data format for each column in the corresponding drop-down menu.

The data are saved in the default data format or in the format set in data assignment (page 38) (Assign data 2/2 dialog).
 - e. Click **OK**.

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.

By default, all values beginning with =, +, -, or @ are automatically masked by a preceding single quotation mark before exported as CSV file. For details, refer to <https://www.contextis.com//resources/blog/comma-separated-vulnerabilities>. You can disable the automatic masking (page 195) in the server configuration.

Note that when exporting data from a widget, the export may fail if the stream size is too large. In this case, you can increase the **maxPostSize** value of the HTTP Connector of your **Tomcat** installation. If not already specified, you can set this attribute to 2097152 (2 megabytes). You can set this value in the **server.xml** file.

5.1.9 Set the step width of a distribution chart

You can change the preset step width of a distribution chart.

Procedure

1. Open a dashboard in dashboard view mode. (page 5)
2. Click a **Distribution chart** widget.
3. Click the  **Menu** icon.
The menu is available if the **Widget menu** option is enabled for the widget.
4. Click **Set step width**.
5. Specify the step width for the distributed measure in the **Step width** input box.
6. Click **OK**.

Your settings are applied.

5.1.10 Use bookmarks

You can use bookmarks to save the dashboard settings that you specified in view mode, such as selections and filter settings.

A bookmark does not store the values from the data sources assigned to the widgets in the dashboard, but the configuration data of an analysis as well as the dashboard settings.

The bookmarks are stored in the separate bookmark bar. Using the bookmark bar, you can create (page 12), share (page 14), rename (page 12), update (page 13) and delete (page 13) bookmarks.

The bookmark bar is hidden by default and must be enabled.

5.1.10.1 Open the bookmark bar

You can open the bookmark bar in your dashboard.

The bookmark bar is hidden by default.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open a dashboard in dashboard view mode. (page 5)
2. Click **Bookmarks** on the right side of the dashboard.

The bookmark bar opens.

To hide the bookmark bar, click **» Hide bookmark bar**.

5.1.10.2 Create bookmarks

You can create bookmarks to save your dashboard settings that you made in view mode, for example, your filter settings.

The bookmarks are stored in the a separate bookmark bar, which is hidden by default.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open a dashboard in dashboard view mode. (page 5)
2. Make your settings, for example, set the filter values for a widget.
3. Open the bookmark bar (page 11).
4. Click **Add bookmark**.

The bookmark is created and added to the bookmark bar.

By default, a new bookmark is marked as private. That means that only the creator of the bookmark can see it in the bookmark bar. You can, however, share your bookmarks (page 14) with any other user who has access to the dashboard.

5.1.10.3 Open a bookmarked dashboard

You can apply the settings you have saved as a bookmark (page 11) to the dashboard.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmark bar (page 11).
2. Click a bookmark.

The Dashboard applies the settings stored in the bookmark.

5.1.10.4 Rename a bookmark

You can rename an existing bookmark (page 11).

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmark bar (page 11).
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Rename**.
4. Enter a name for the bookmark.
5. Press the **Enter** key.

The selected bookmark is renamed.

5.1.10.5 Delete a bookmark

You can delete an existing bookmark (page 11).

Warning

A deleted bookmark cannot be restored.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmark bar (page 11).
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Delete**.
4. Click **Yes**.

The selected bookmark is deleted.

5.1.10.6 Update a bookmark

You can update an existing bookmark with the current settings of your displayed dashboard. The saved settings of the bookmark will be overwritten.

Warning

An updated bookmark cannot be restored.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmark bar (page 11).
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Update**.
4. Click **Yes**.

The saved dashboard settings of the selected bookmark are updated.

5.1.10.7 Share a bookmark

You can share your bookmarks with any other user who has access to the dashboard.

You can mark your bookmarks as follows:

- **Private** (default setting): Only the creator of the bookmark can see it in the bookmark bar.
- **Public**: Anyone who has access to the dashboard can see the bookmark.
- **Shared**: Anyone who has access to the dashboard can use the bookmark using a specific URL. You can copy the bookmark URL to the clipboard and share it with other users.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmark bar (page 11).
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Share**.
4. Make your settings.

Your settings are applied.

5.2 Use dashboard in edit mode

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

To use a dashboard in edit mode, you need the **Edit** permission (page 20) for the dashboard.

5.2.1 Open a dashboard in the dashboard editor

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

Prerequisite

You have the required access privileges for the dashboard (page 20).

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.

Warning

Note that when you save your dashboard using the name of an existing dashboard, conflicts may occur when you reopen the dashboard.

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 (page 20) 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 14) 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 37).

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 37).
2. Resize the layout group. See Resize widgets (page 82).
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 14) or open a dashboard (page 19) 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 14).
2. Click **Manage** > **Switch to fixed grid** in the dashboard main menu.

The fixed-grid workspace is displayed.

5.2.4 Save a dashboard

You can save a dashboard and give it a unique name, for example, to make a copy of your dashboard.

Warning

Note that when you save your dashboard using the name of an existing dashboard, conflicts may occur when you reopen the dashboard.

Procedure

1. Click **Manage** > **Save** in the main menu of the dashboard editor.
You can create a copy of the currently opened dashboard by using the **Save as** option.
2. Specify your settings.

The dashboard is saved on the server.

You can change your settings by editing the dashboard (page 19).

5.2.5 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 15) 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.

Warning

Note that when you save your dashboard using the name of an existing dashboard, conflicts may occur when you reopen the dashboard.

Your changes are applied.

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

5.2.6 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 19)
2. Click **Manage > Delete** in the dashboard main menu.
3. Click **Yes**.

The selected dashboard is deleted.

5.2.7 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 14) or open a dashboard (page 19) 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.

5.2.8 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 14) or open a dashboard (page 19) 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. 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.

5. Enable or disable the **View** or **Edit** permissions of a user or a user group.
6. 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.

7. 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.

8. 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.

5.2.9 Change the dashboard style

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

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) in the dashboard editor.
2. Click **Manage > Change style template** in the dashboard main menu.
3. Select a style in the **Dashboard** drop-down menu.

4. Click **OK**.

The selected style is applied to the current dashboard.

5.2.10 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.

5.2.10.1 Create a Process Mining context

You can create a Process Mining context (page 22) 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 14) or open a dashboard (page 19) 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 from the **Language** drop-down menu.

Only the languages specified for the PPM client of the selected PPM connection are available.

7. Select the processes to be analyzed from 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 for measures and dimensions** drop-down menu.

The menu provides the relations specified for the PPM client of the selected PPM connection.

Based on the relation you selected, the corresponding measure and dimension are provided for the context-based widgets, except for the **Function Flow** widget (see step 9).

For the Conformance check widget select Process and relation: 'Conformance issue'.

By default, Process and Function is selected.

Note that if you are using PPM in a version older than 10.5, the drop-down menu is not available.

9. Select the relation to be analyzed from the **Relation for Function Flow widget** drop-down menu.

The menu provides the relations available for the **Function Flow** widget that are specified for the PPM client of the selected PPM connection. Depending on the selected process type, you can choose between different relations.

Based on the relation your selected, the corresponding measure and dimension are provided for the context-based **Function Flow** widget.

By default, the **Function follows** relation is selected. If this relation is not available, the first suitable relation is selected. If no suitable relation is available for the selected process type, None is selected.

10. Select the measures and dimensions to be analyzed.

The measures and dimensions provided depend on the process type and relation 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.

11. 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.

12. Click **Next** to manually edit the data columns of the context.

13. Enter a term in the **Search** field to filter the data columns list.

14. 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.

15. 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.

16. 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**.
17. Click the **Delete** icon to delete a data column from the Process Mining context.
18. 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.

5.2.10.2 Edit a Process Mining context

You can change the settings of the Process Mining context (page 22) 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 30).

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 22). 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.

5.2.10.3 Delete a Process Mining context

You can delete a Process Mining context (page 22) configured for a dashboard.

Warning

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 14).
2. Click **Options** > **Process Mining** > **Delete context** in the dashboard main menu.
3. Click **Yes**.

The Process Mining context is deleted.

5.2.10.4 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 22). The procedure to assign data columns is similar to the standard method for assigning data columns to widgets (page 38) (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 22).

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 (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 (page 38) 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 (page 38). If you assign another data source, you cannot reassign the Process Mining context to the widget.

5.2.10.5 Configure filters for Process Mining context-based widgets

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

By default, the filter conditions are also automatically set when you assign data columns of the Process Mining context to the widgets (page 25). 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 22).

Procedure

1. Open a dashboard in the dashboard editor (page 14).
2. Assign data columns to Process Mining context-based widgets (page 25).
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. 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 29). 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).

5.2.10.5.1 Configure a date filter for a Process Mining context

Using the **Date filter** (page 279) widget, you can specify a filter column for filtering the widgets of a Process Mining context-based dashboard. You can select a context column of **Date** type to specify a single date value, such as 2018-03-25, or a time range with a start value and an end value, for example, from 2018-03 to 2018-07.

Note that if you are using a PPM version older than 10.5, the **Allow range selection** option is not available for Process Mining context-based dashboards.

Procedure

1. Click a **Date filter** widget on the dashboard. The relevant properties dialog is displayed.

2. To specify a time range, you must first enable the **Allow range selection** option on the **Configure** tab.

Do not specify the **From** and **To** parameters. They are set automatically by the context. If only one range value (the start value or the end value) is set in the context, the filter is interpreted as open range, for example, > 2010 or < 2010.

3. Click **Context** in the properties dialog to open the corresponding tab.
4. 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.

Note that depending on your selection on the **Context** tab, different entries are available, if you specify an action for a widget (page 76), for example, Set selection (page 77) or Call URL (page 77).

- If you select **None** on the **Context** tab, the available entries are **from** and **to**.
- If you select a column, for example, 'Time [By month]', and the **Allow range selection** option is not enabled, the available entries are 'Time [By month]_from' and 'to'.
- If you select a column, for example, 'Time [By month]', and the **Allow range selection** option is enabled, the available entries are 'Time [By month]_from' and 'Time [By month]_to'.

5.2.10.6 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 22).

Procedure

1. Open a dashboard in the dashboard editor (page 14).
2. Assign data columns to Process Mining context-based widgets (page 25).
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.

5.2.10.7 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 22).

Procedure

1. Open a dashboard in the dashboard editor (page 14).
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 22).

The filter panel is enabled and available in view mode.

You can configure the filter panel in dashboard edit mode. (page 29)

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.

5.2.10.8 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 296**) widget for details.

Prerequisites

The filter panel is enabled (page 29).

Procedure

1. Open a dashboard in the dashboard editor (page 14).
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 25).
 - 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 26) 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)

5.2.10.9 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 24).
3. Click **OK** to save your changes.
4. Software AG recommends saving your dashboard as a copy to keep the settings of your old dashboard.
 - a. Click **Manage > Save as** in the dashboard main menu.
 - b. Make your settings.
 - c. Click **Save**.
5. Click a 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 25)

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.

5.2.11 Use Multi-lingual dashboards

With version 10.3, MashZone NextGen supports multi-lingual dashboards. Every widget can offer different elements for translation. For example, this includes 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).

5.2.11.1 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 14) or open a dashboard (page 19) 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 from 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 32)

5.2.11.2 Translate the widget content

You can translate the widget content in all languages supported by a dashboard (page 31). 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. (page 38) Widgets to which no data can be assigned are excluded, such as the **Label** widget.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 (page 38), 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.

5.2.11.3 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 31) and the widget contents are translated (page 32).

Procedure

1. Display a dashboard in dashboard view mode. (page 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 31)
The dashboard is displayed in the language selected.

5.2.11.4 Translate the tab title

You can translate the tab titles in all languages supported by a dashboard (page 31).

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) in the dashboard editor.
2. Add one or more tabs to the dashboard. (page 84)
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.

5.2.12 Manage bookmarks

In Edit mode, you can manage the created bookmarks of a dashboard (page 11).

The bookmarks are stored in the separate bookmark bar. Using the bookmark bar, you can rename (page 34), share (page 35), and delete (page 34) bookmarks, or set a bookmark as default (page 35).

The bookmark bar is hidden by default and must be enabled for it to be used.

5.2.12.1 Enable the bookmark bar

You can enable the bookmark bar in Edit mode. If you enable the bookmark bar in the dashboard editor, the bookmark bar is also enabled in View mode.

The bookmark bar is hidden by default.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) in the dashboard editor.
2. Click **Options** > **Bookmarks** in the main menu.

The bookmarks bar is enabled.

To open the bookmark bar (page 11), click **Bookmarks** on the right side of the dashboard.

5.2.12.2 Open the bookmark bar

You can open the bookmark bar in Edit mode.

The bookmark bar is hidden by default.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) in the dashboard editor.
2. Click **Bookmarks** on the right side of the dashboard.

The bookmark bar opens.

To hide the bookmark bar click **»» Hide bookmark bar**.

5.2.12.3 Rename a bookmark

You can rename an existing bookmark (page 11) in Edit mode.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmarks bar. (page 34)
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Rename**.
4. Enter a name for the bookmark.
5. Press the **Enter** key.

The selected bookmark is renamed.

5.2.12.4 Delete a bookmark

You can delete an existing bookmark (page 11) in Edit mode.

Warning

A deleted bookmark cannot be restored.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmarks bar. (page 34)
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Delete**.
4. Click **Yes**.

The selected bookmark is deleted.

5.2.12.5 Share a bookmark

You can share your bookmarks with any other user who has access to the dashboard.

You can mark your bookmarks as follows:

- **Private** (default setting): Only the creator of the bookmark can see it in the bookmarks bar.
- **Public**: Everyone who has access to the dashboard can see the bookmark.
- **Shared**: Everyone who has access to the dashboard can use the bookmark using a specific URL. You can copy the bookmark URL to the clipboard and share it with other users.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmarks bar. (page 34)
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Share**.
4. Make your settings.

Your settings are applied.

5.2.12.6 Set and reset a bookmark as default

You can set or reset a bookmark as the default bookmark to load when you open the dashboard in View mode. The bookmark settings are automatically applied to the dashboard.

Prerequisite

The bookmark bar is enabled. (page 33)

Procedure

1. Open the bookmark bar. (page 34)
2. Move your mouse pointer over a bookmark and click the  menu icon.
3. Click **Set bookmark as default**.
4. To reset the default bookmark, move the mouse pointer over a default bookmark, click the  menu icon and then **Reset default bookmark**.

Your settings are applied.

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 273) and the relevant configurable parameters:

-  Line chart (page 289)
-  Column chart (page 277)
-  Bar chart (page 273)
-  Pie chart (page 293)
-  Bubble chart (page 275)
-  Grid (page 283)
-  Circular gauge chart (page 276)
-  Horizontal and vertical gauge chart (page 285)
-  Traffic light (page 300)
-  Drop-down box (page 282)
-  Input field (page 287)
-  Image (page 286)
-  Label (page 288)
-  Function Flow diagram
-  Jump to PPM (page 299)

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 (page 38) for details.
4. Click  **Configure filter** to specify filter conditions. See Define filters for widgets (page 74) for details.
5. Click the **Config** tab and set the widget display options. See the list of widgets (page 273) 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 85) 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 14) or open a dashboard (page 19) 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 (page 38) for details.
4. Click the  **Configure filter** icon to specify filter conditions. See Define filters for widgets (page 74) for details.
5. Click the **Config** tab and set the widget display options. See the list of widget (page 273) 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 85) 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 70). Data sources are not required for **Input field** and **Image** widgets.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 70).
5. Select a data source operator in the **Add a data source** bar and specify your settings. See Data source operators (page 303) 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 323) for a list of available data transformation operators.
7. To calculate a data preview of an operator, click the  **Calculate preview** icon. This allows you to track all data changes step by step. For details, see chapter Calculate the feed data (page 71).
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 289), column chart (page 277), or bar chart (page 273), 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 (page 38).

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 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 (page 74) 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 42) for details.
7. Click **Coloring** and then click the **Rating** tab to specify the KPI rating. See Configure KPI rating (page 42) 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 position	Specifies on which side of the chart the axis is located. Default is left.
Scale type	Specifies the scale type of the axis values. The scale type is available for axes with numeric values. <ul style="list-style-type: none"> ▪ Linear: The axis values are displayed linearly. This is the default scale type. ▪ Log (logarithmic): The maximum value of the source column is set to 100 and the minimum is set to 1. All axis values are between 100 and 1. Negative values and 0 values of the source column are not taken into account.
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 KPIs, data points, or tool tips. By default, the data source column name is used.

Option	Description
Format	<p>Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type.</p> <p>For details on setting the date format, refer to the chapter <i>Change data type - single value</i> (page 333).</p> <p>Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week, you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.</p>
Sorting	<p>Sorts the labels of the X-axis data points in ascending or descending order, or sorts the labels of the X-axis data points 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.</p>
Sort by	<p>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.</p>
Round numerically	<p>Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.</p>
Aggregation	<p>Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.</p> <p>The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.</p>
Text tab	<p>Enables you to specify the Axis title and Axis format.</p>
Data range tab	<p>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.</p> <p>You can enter specific values manually or assign data source columns for a dynamic value assignment using drag and drop.</p>
Coloring	<p>Enables you to specify KPI thresholds and ratings. You can set the background or the graphic color depending on the KPI values.</p>
Thresholds tab	<p>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.</p>

Option	Description
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 a distribution chart

To display data in a distribution chart (page 281), you must assign the relevant data source columns to the required widget elements.

The distribution chart is available only for context based dashboards.

Prerequisite

You have created a Process Mining context for the dashboard (page 22).

Procedure

1. Click a **Distribution chart** 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 (2/3)** dialog is displayed.
3. Select a measure to be distributed on the X-axis from the **Measure to distribute** drop-down menu.
4. Specify the step width for the distributed measure in the **Step width** input box.
In view mode, you can change the step width (page 11) in the settings dialog of the distribution chart.
5. If you want to add or change the data columns provided by the Process Mining context, click **Edit context**.
6. Click **Next**.

The **Assign data (3/3)** dialog is displayed. Here you can assign the data source columns to the widget elements. The data column assigned to the X-axis is preset by your selection in the **Assign data (2/3)** dialog and cannot be changed.

7. Assign a data column to the **Y-axis** using drag and drop.
8. 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 only for filtering widgets. See Define filters for widgets (page 74) for details.

The data source columns are assigned to the widget elements.

5.3.4.2.1 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 **X-axis** box to specify the X-axis settings.
2. Click the **Y-axis** box to specify the Y-axis settings.
 - a. Click the **General** tab to specify the position, title, and format of the axis.
 - 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 a source column assigned to the **More columns (invisible)** box and specify the settings.
5. 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 position	Specifies on which side of the chart the axis is located. Default is left.
Scale type	Specifies the scale type of the axis values. The scale type is available for axes with numeric values. <ul style="list-style-type: none"> ▪ Linear: The axis values are displayed linearly. This is the default scale type. ▪ Log (logarithmic): The maximum value of the source column is set to 100 and the minimum is set to 1. All axis values are between 100 and 1. Negative values and 0 values of the source column are not taken into account.

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 KPIs, data points, or tool tips. By default, the data source column name is used.
Format	<p>Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type.</p> <p>For details on setting the date format, refer to the chapter Change data type - single value (page 333).</p> <p>Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week, you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.</p>
Sorting	Sorts the labels of the X-axis data points in ascending or descending order, or sorts the labels of the X-axis data points 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	<p>Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.</p> <p>The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.</p>

5.3.4.3 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 293) can display one numeric KPI iterated via a dimension (text or date dimension).

Prerequisite

You have assigned a data source to the widget (page 38).

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 **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 (page 74) 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, 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 tool tips. Available for columns of date and numeric type.

Option	Description
	<p>For details on setting the date format, refer to the chapter Change data type - single value (page 333).</p> <p>Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week, you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.</p>
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	<p>Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.</p> <p>The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.</p>

5.3.4.4 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 275) 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 (page 38).

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 (page 74) 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 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 49) 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
Scale type	Specifies the scale type of the axis values. The scale type is available for axes with numeric values. <ul style="list-style-type: none"> ▪ Linear: The axis values are displayed linearly. This is the default scale type. ▪ Log (logarithmic): The maximum value of the source column is set to 100 and the minimum is set to 1. All axis values are between 100 and 1. Negative values and 0 values of the source column are not taken into account.
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.

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	<p>Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type.</p> <p>For details on setting the date format, refer to the chapter <i>Change data type - single value</i> (page 333).</p> <p>Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week, you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.</p>
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	<p>Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.</p> <p>The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.</p>
Text tab	Allows you to specify the Axis title and Axis format .
Data range tab	<p>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.</p> <p>You can enter specific values manually or assign data source columns for a dynamic value assignment using drag and drop.</p>
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.

5.3.4.4.2 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.

5.3.4.5 Assign data columns to grids

To display data in the Grid (page 283) 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 (page 38).

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 **Grid columns** 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 (page 74) 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 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 52) 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.
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, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
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.

Option	Description
Set as key column for selection	<p>Sets the selected column as key column.</p> <p>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.</p> <p>You can select one or more columns as key columns.</p>
Text tab	Allows you to specify the New column name, Format and Aggregation.
Thresholds tab	Allows you to specify the thresholds values range.

5.3.4.5.2 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.

5.3.4.6 Assign data columns to lists

To display data in the List (page 290) widget, you must assign the relevant data source columns to the widget elements required.

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

Prerequisite

You have assigned a data source to the widget (page 38).

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 (page 74) for details.

The data source columns are assigned to the widget elements.

5.3.4.6.1 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	<p>Output format of the column values, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333).</p> <p>Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week, you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.</p>

Option	Description
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.

5.3.4.7 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 276) 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 (page 38).

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 56) 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, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
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.

5.3.4.7.2 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.

- 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.

5.3.4.8 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 285) 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 (page 38).

Procedure

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

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

5.3.4.8.1 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.

- Click the source column assigned to the **KPI** and specify the settings.
- Click **Thresholds** to specify the KPI thresholds. See Configure thresholds (page 58) for details.
- 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, for example, used for data points or tool tips. Available for columns of date and numeric type.

Option	Description
	<p>For details on setting the date format, refer to the chapter Change data type - single value (page 333).</p> <p>Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week, you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.</p>
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Aggregation	<p>Specifies how the KPI value is calculated. Available for numeric columns. Default is Average value.</p> <p>The number of aggregation types provided by context-based widgets may differ from that of non-context-based widgets.</p>
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.

5.3.4.8.2 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.

5.3.4.9 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 300) 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 (page 38).

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.

5.3.4.9.1 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 60) 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, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd

Option	Description
	instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
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.

5.3.4.9.2 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.

5.3.4.10 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 282) 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 (page 38).

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 (page 74) for details.

The data source columns are assigned to the widget elements.

5.3.4.10.1 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, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For

Option	Description
	the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
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.

5.3.4.11 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 288) 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 (page 38).

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.

5.3.4.11.1 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, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.
Initial selected value	Specifies the initial selected column value. By default, the first column value is selected.
Sort descending	Displays the column values in a descending order in the drop-down menu.

5.3.4.12 Assign data columns to a rich text area

To display data in the Rich text area (page 294) widget, you must assign the relevant data source columns to the required widget elements.

Prerequisite

You have assigned a data source to the widget (page 38).

Procedure

1. Click a Rich text area (page 294) 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.

5.3.4.12.1 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, for example, used for data points or tool tips. Available for columns of date and numeric type. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
Round numerically	Displays rounded KPI values. Available for numeric columns. Enabled by default.

5.3.4.13 Assign data columns to sliders

5.3.4.13.1 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.

5.3.4.14 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 279) 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 279) for details.

Prerequisite

You have assigned a data source to the widget (page 38).

Procedure

1. Click a Date filter (page 279) 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.

5.3.4.15 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 301) 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 (page 38).

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.

5.3.4.15.1 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 **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 67) 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 will be selected by default.

Option	Description
Text tab	Enables you to specify the column settings.
Thresholds tab	Enables you to specify the threshold values range and to set the background and graphic color.
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. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
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 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 292) 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 (page 38).

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.

5.3.4.16.1 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 69) 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. For details on setting the date format, refer to the chapter Change data type - single value (page 333). Note that the pattern for the day of the week used in the Assign Data 2/2 dialog is different from the pattern used in data feed processing. For the date format of the day of the week , you must use the pattern dddd instead of eeee as in the data feed processing. For example, the format dddd.MM.YYYY returns Monday.12.2019.
Round numerically	Displays rounded KPI values in the chart. Available for numeric columns. Enabled by default.

5.3.4.16.2 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.

5.3.4.17 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 (page 38).
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 (page 74) for details.

5.3.4.18 Calculate the feed data

You can calculate the data for all operators of the feed definition and display the corresponding preview in a preview table. This enables you to track all data changes step by step.

The preview table provides several preview modes. For details, see the table below.

For performance reasons, the preview is limited to a maximum of 1000 rows by default.

Procedure

1. Click the  Calculate preview icon in the header of an operator.

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

2. Select the preview mode in the corresponding drop-down menu in the header of the preview table.

The raw data preview (quick or full) of text-file operators (XML, JSON, CSV) is displayed in a single (non-sortable) column. The raw preview of an **Excel** operator has Excel column names (A, B, C, etc.) and an initial column that contains the line number.

3. To change the sort order of a table, click a column heading.

Click once to change the sort order to ascending.

Click twice to change the sort order to descending.

Click three times to change the sort order to default.

The result is displayed according to your settings.

The preview mode remains selected as long as the data feed (or **Assign data 1/2** (page 38) dialog) is open.

The preview table provides the following preview modes.

Preview mode	Description
Quick	<p>The selected operator returns at most 1000 rows in its result. This applies to the operator for which a preview is needed, as well as to all other operators required for calculating the preview.</p> <p>This preview mode is equal to the data preview provided in MashZone NextGen up to version 10.3.</p> <p>The Quick preview mode is preselected by default.</p>
Quick raw	<p>The preview displays the data of the source of the selected operator.</p> <p>This preview mode is supported only for the data source operators XML, JSON, CSV, and Excel. XML and JSON source data are reformatted with indentation. The results are limited to 1000 rows.</p>
Full	<p>The preview is displayed without limiting the operator results to 1000 rows. The actual limit depends on the license you use.</p>
Full raw	<p>The result of the raw data preview is displayed without limiting it to 1000 rows. The actual limit depends on the license you use.</p>
Statistics	<p>The result table has one line per measure, an initial column containing the name of the measure, and one column for each column of the operator's result table.</p> <ul style="list-style-type: none"> ▪ Count: the count of non-NULL values for this column ▪ #Null: the count of NULL values for this column ▪ Min: the minimum value of this column ▪ Max: the maximum value of this column ▪ Avg: the average value of this column ▪ Std: the standard deviation of this column ▪ Unique: the number of unique values in this column ▪ Top: a value that has the highest frequency ▪ Freq: the frequency of the top value

5.3.4.19 Specify the color mapping for dimensions

You can specify a color mapping to visualize the values of a text dimension in a widget. By default, the color mapping is applied to the column assigned to the **Partition** element of the widget. If there is not partition set, the color mapping is automatically applied to the X-axis. You can use an existing column in the data feed that provides the required color mapping for the dimension. The corresponding column must contain the colors in hexadecimal format (for example, #3482B7), which is accepted by the browser.

By default, the colors of the data points in a chart (such as columns or bubbles) are provided by the color set of the theme applied to the dashboard or the style assigned to the widget. If

you specify a color mapping, the mapping is preferred to the default settings. The colors specified for thresholds and ratings are preferred to the color mapping.

The color mapping is available for **Column**, **Bar**, **Line**, **Bubble**, and **Pie** chart.

Prerequisite

You have assigned a data source to the widget (page 38).

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. If you have not already done so, assign columns to the widget elements using drag & drop.
4. Click **Coloring -> Mapping** to open the corresponding tab.
5. Enable the **Data column** option to use a data column for the column mapping.
The **Default** option is enabled by default. This means that the color specified by the dashboard theme or widget style is used.
6. Assign a data column containing the color mapping to the **Data column** box using drag and drop.
7. Click **OK**.

The color mapping is applied.

Tip

You can also change the coloring of the data points when you configure thresholds and ratings for a measure.

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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 (page 5)).

Most widgets support data filters. Especially, the **Drop-down box** (page 282), **Input field (page 287)** and **List (page 290)** widgets can be used for filtering other widgets. (See Insert widgets to a dashboard (page 37)).

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 (page 5) for details.

In particular, the List (page 290) 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 70)) of **List** type in the data feed, Date user input (List) (page 361), Number user input (List) (page 363), and Text user input (List) (page 364).

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 26) 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 14) or open a dashboard (page 19) 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 82) 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 (page 38) 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 70) 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 (page 38) 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  **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 (page 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 another 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 50) for details.

You can use the defined action in the dashboard view mode. See Open a dashboard in MashZone NextGen (page 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 84).

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

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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. Enable 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 that you want to select by your action. For this, drag the relevant coordinate and drop it in the field in the **Select selection component** 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 14) or open a dashboard (page 19) 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 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 286) 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 14) or open a dashboard (page 19) 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 an 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 201) 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 at 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.

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 85)), 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 14) or open a dashboard (page 19) 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 (page 261) for details.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 \(page 18\)](#) for details.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 82).

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 84) 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 the widget style

You can assign a different style to a widget on your dashboard. Using styles you can customize the look and feel of your widgets, for example, colors schemes, fonts or background color.

The styles provided for the selected widget are part of the style template that is assigned to the current dashboard (page 21).

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 in the **Style** drop-down menu.

The selected style 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 14) or open a dashboard (page 19) 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 (page 38).
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-  
4cb1-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 14) or open a dashboard (page 19) 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 currently 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 18) for details.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 18) for details.

Procedure

1. Create a dashboard (page 14) or open a dashboard (page 19) 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 18).

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 20).

- a. Open the **Manage dashboard permission** dialog for the dashboard.
- b. Add the **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 **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, 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 assign the **Guest** user group to the relevant file resource.

The **Guest** user group only has the **View** permission.

Procedure

1. Grant file permissions to the **Guest** user group.
 - a. Open the MashZone NextGen Admin Console (page 106).
 - b. Click **Platform Features > File resources**.
 - c. Click  **Edit File Permissions**.
 - d. Assign the **Guest** user to the principals with permissions.
2. 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 286) widget.

5.5.6 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 14) or open a dashboard (page 19) 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
```

5.5.7 Export and import dashboards

You can export and import dashboards using the MashZone NextGen Admin Console. For details on exporting and importing dashboards, see the chapter Deploying artifacts using the Admin Console (page 212).

5.5.8 Customize styles and themes

You can customize the styles and themes that are used to specify the appearance of the application, dashboards, and widgets.

For example, you can replace the logo in the application header, change the text displayed on the welcome page, create your own dashboard and widget style templates, or adjust the font size to different devices.

For details on how to customize styles and themes, see the chapter [Manage application and widget styles](#) (page 226).

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 96) 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 91)
- Select data sources (page 91)
- Calculate the feed data (page 71)
- Add further operators (page 93)
- Connect the operators (page 94)
- Save the data feed (page 94)

6.1.1 Open the data feed editor

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 303) 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 a preview table. This enables you to track all data changes step by step.

The preview table provides several preview modes. For details, see the table below.

For performance reasons, the preview is limited to a maximum of 1000 rows by default.

Procedure

1. Click the  Calculate preview icon in the header of an operator.
The feed data is calculated for the selected operator of the feed definition. The result is displayed in a preview table at the lower edge of the workspace.
2. Select the preview mode in the corresponding drop-down menu in the header of the preview table.
The raw data preview (quick or full) of text-file operators (XML, JSON, CSV) is displayed in a single (non-sortable) column. The raw preview of an **Excel** operator has Excel column names (A, B, C, etc.) and an initial column that contains the line number.
3. To change the sort order of a table, click a column heading.
Click once to change the sort order to ascending.
Click twice to change the sort order to descending.
Click three times to change the sort order to default.

The result is displayed according to your settings.

The preview mode remains selected as long as the data feed (or **Assign data 1/2** (page 38) dialog) is open.

The preview table provides the following preview modes.

Preview mode	Description
Quick	<p>The selected operator returns at most 1000 rows in its result. This applies to the operator for which a preview is needed, as well as to all other operators required for calculating the preview.</p> <p>This preview mode is equal to the data preview provided in MashZone NextGen up to version 10.3.</p> <p>The Quick preview mode is preselected by default.</p>
Quick raw	<p>The preview displays the data of the source of the selected operator.</p> <p>This preview mode is supported only for the data source operators XML, JSON, CSV, and Excel. XML and JSON source data are reformatted with indentation. The results are limited to 1000 rows.</p>
Full	<p>The preview is displayed without limiting the operator results to 1000 rows. The actual limit depends on the license you use.</p>
Full raw	<p>The result of the raw data preview is displayed without limiting it to 1000 rows. The actual limit depends on the license you use.</p>
Statistics	<p>The result table has one line per measure, an initial column containing the name of the measure, and one column for each column of the operator's result table.</p> <ul style="list-style-type: none"> ▪ Count: the count of non-NULL values for this column ▪ #Null: the count of NULL values for this column ▪ Min: the minimum value of this column ▪ Max: the maximum value of this column ▪ Avg: the average value of this column ▪ Std: the standard deviation of this column ▪ Unique: the number of unique values in this column ▪ Top: a value that has the highest frequency ▪ Freq: the frequency of the top value

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 323) 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.

Warning

Note that when you save your data feed using the name of an existing data feed, conflicts may occur when you reopen the data feed.

Procedure

1. Click **Manage > Save** in the main menu.
You can create a copy of the currently opened data feed using 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 95).

6.2 Edit a data feed

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

Prerequisite

You have the required access privileges for the data feed (page 96).

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.

Warning

Note that when you save your data feed using the name of an existing data feed, conflicts may occur when you reopen the data feed.

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 14) or open a dashboard (page 19) 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 (page 91) or open an existing data feed (page 95).
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 (page 91) or open an existing data feed (page 95).
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. 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.
5. Enable or disable the **View** or **Edit** privileges of a user or a user group.
6. 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.

7. 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.

8. 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 Data feed operators

The data feed editor provides a wide range of data source and data transformation operators for creating data feed definitions (page 91).

- Data source operators (page 303)
- Data transformation operators (page 323)
- User input operators (page 360)

6.7 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 of those specific results. For that purpose, you can materialize the results of a data feed evaluation in Terracotta, 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.

Working with Materialized Feeds affects several components and workflows:

- You must configure a Terracotta connection that enables you to create a dataset on the Terracotta server and materialize feed results in it.
For details on configuring a Terracotta connection, see [Manage Terracotta connections](#).
- You must configure a Materialized Feed that contains the definition of which data feed should be materialized in which dataset using which Terracotta connection.
For details on configuring a Materialized Feed, see [Manage Materialized Feeds \(page 190\)](#).
- You can either use the Admin Console, the MashZone NextGen administration tool, or publicly available REST web services in order to calculate data feed results and materialize them subsequently in Terracotta.
For details on triggering a Materialized Feed, see [Materialize data feed results in Terracotta \(page 193\)](#).
- You can access and analyze data feed results materialized in Terracotta using the Terracotta source operator (page 317).
- 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 192).

6.7.1 Use datasets storing data feed results

To access a dataset with Materialized Feed results on a Terracotta server, you must first configure a Terracotta connection and a Materialized Feed in MashZone NextGen. You can define both in the MashZone NextGen Admin console. See Manage Terracotta connections and Manage Materialized Feeds (page 190).

Procedure

1. Create a new data feed (page 91) or open an existing data feed (page 95).
2. Insert a **Terracotta** (page 317) data source. For details, see Select data sources (page 91).
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 the **Dataset Name** box.
5. Configure the RAQL query, refresh rate, and input parameters. For details, see **Terracotta** (page 317) data source.

Your settings are applied.

6.7.2 Manage Materialized Feeds in Terracotta

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

- A Materialized Feed is deleted (page 191) 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 191) and the name of the dataset is changed, the old dataset will not be deleted.

If you want to delete these datasets, a Terracotta administrator must perform that task using Terracotta 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 not be manipulated or deleted because it is required for the correct interaction of MashZone NextGen with Materialized Feed datasets.

6.7.3 Use datasets from an external application

To access a dataset generated by an external application on a Terracotta server, you must first configure a Terracotta connection in MashZone NextGen. You can define Terracotta connections in the MashZone NextGen Admin console. For details, see [Manage Terracotta connections](#).

Procedure

1. Create a new data feed (page 91) or open an existing data feed (page 95).
2. Insert a **Terracotta** (page 317) data source. For details, see [Select data sources](#) (page 91).
3. Select **Dataset** in the **Source** drop-down menu.
4. Select a Terracotta connection in the **Terracotta alias** drop-down menu. The alias specifies the Terracotta server with which you want to connect. The menu contains all available Terracotta connections for which you have **Usage** privilege.
5. Select a dataset alias in the **Dataset name** drop-down menu. The menu contains all available Terracotta datasets for the selected Terracotta alias. For details, see [Manage Terracotta connections](#). Datasets generated by MashZone NextGen for Materialized Feeds on the Terracotta server are excluded from the list.
6. Click **Configure columns** to configure the columns of the dataset.

When you select 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. Disable 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 match the cell definition in the store. If does not fit the definition, the corresponding values are always null. For example, if the cell definition is age with type INT and you specify age with type DOUBLE, the column definition does not correspond to the cell definition and the values are null.

7. Configure the RAQL query, refresh rate, and input parameters. For details, see **Terracotta** (page 317) data source.
8. Your settings are applied.

6.7.4 Which privileges are required?

In addition to configuring the sharing settings for Materialized Feeds (page 192), you must also configure the required privileges for the data feed and the Terracotta connection used.

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

- In dashboard mode the contents of a dataset with Materialized Feed results.
 - Data feed accessing the dataset with Materialized Feed results: Enable the **Display** privilege (page 96).
 - Materialized Feed: Enable the **Display** privilege (page 192).
 - Terracotta connection: Enable the **Display** privilege.
- In the feed editor, specify the Materialized Feed alias in the list of available aliases in the Terracotta source operator.
 - Materialized Feed: Enable the **Usage** privilege (page 192).
 - Terracotta connection: Enable the **Usage** privilege.
- Calculate a data feed and materialize its results in a dataset being managed by a Terracotta connection. Only users with administrator privileges, who have the following artifact rights can do this.
 - Data feed whose results are materialized: Enable the **Display** privilege (page 96).
 - Materialized Feed: Enable the **Usage** privilege (page 192).
 - Terracotta connection: Enable the **Usage** privilege.

6.8 Other

6.8.1 Export and import data feeds

You can export and import data feeds using the MashZone NextGen Admin Console.

For details on exporting and importing data feeds, see the chapter Deploying artifacts using the Admin Console (page 212).

6.8.2 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 (page 91) or open an existing data feed (page 95).
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 Good to know

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

7.1 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 them.

Dashboards make it easy to visualize and analyze information. You can combine data from any original source and visualize them using 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 14) for instructions.

7.2 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 (page 38) 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 (page 91) for instructions.

Data feed definition and feed table



7.3 What is a 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 a view-only license.

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

- The user can view dashboards in only view mode.
- 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 or she can Manage dashboard permissions (page 20) and/or Manage data feed permissions (page 96).
- 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 149).

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 enables you to control user interactions including registering or creating dashboards and data feeds. You can also secure access for all users to work with these artifacts, based on policies that you define.

- **Change password:** For security reason, we strongly recommend that the MashZone NextGen administrator change the standard MashZone NextGen password after installation.
- **Change password of target data sources:** For security reason, we strongly recommend that you change the key 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/. You can change the key by using the **padmin generateKey -a AES -f authTokenKey** command that creates a new **authTokenKey** file. First, you must create a backup of the existing **authTokenKey** file and then copy the new file to that folder. The file should be changed only with an empty repository, because already encrypted passwords can no longer be decrypted. The same applies to exported content. The system where the content is imported must 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 108) for details.
- Incorporate password policies and expiring passwords.

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 using 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 services).
- Hide network ports used solely for internal communication between server widgets.
- Set up secure communication between the client and server using HTTPS. For details, see [Configure HTTPS and Certificate Stores in the Application Server](#) (page 164).
- 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 107), if permitted
- Change your password (page 107), 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. Click the  user icon in the program bar.
2. Click **My Profile**.
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. Click the  user icon in the program bar.
2. Click **My Profile**.
3. Click **My Password**.

4. Enter your new password and confirm this.
5. 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 108) 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, 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 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 110) for more information.
- Valid credentials. See Valid Credentials (page 109) for more information.
- Guest access header or parameter information.

8.2.4.1 User Authentication

MashZone NextGen is initially installed with a set of Default User Accounts (page 110) 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 138), Manage Users (page 139) and Manage User Groups (page 140) 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.
- TLS and User Certificates
- A configurable Single Sign-On solution
- See Authentication with single sign-on solutions (page 110) for configuration instructions.

Permission to work with MashZone NextGen artifacts can also be granted to guests (unauthenticated users), if needed.

8.2.4.2 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 138) for more information.
- TLS 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 138) 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 110) 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 110) for more information.

8.2.4.3 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.

8.2.5 Default user account

MashZone NextGen has one user account that you can use 'out-of-the-box' to access MashZone NextGen dashboards and data feeds.

Username	Password	Built-in Group / Permissions	Description
Administrator	manage	Administrator	A MashZone NextGen administrator.

If you configure MashZone NextGen to use your LDAP Directory as the MashZone NextGen User Repository, the default user account is automatically disabled. If you use the Default User Repository, you can delete this user account in the Admin Console.

Important: You must ensure that at least one user account has MashZone NextGen administrator role.

8.2.6 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 111) 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 113) for details.

8.2.6.1 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 112) for details.

Procedure

1. If needed, configure the MashZone NextGen User Repository. See [Use the Default MashZone NextGen User Repository](#) (page 138) 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 121) for instructions.

8.2.6.2 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.

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 218) 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.

8.2.6.3 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".

8.2.7 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 changing 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 using 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\`

8.2.8 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 prevents clickjacking attacks by 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 [Configure MashZone NextGen server to work with iFrame \(page 158\)](#).

8.2.8.1 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

antiClickJackingUri 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.

8.2.8.1.1 Example

The website <http://website-a.com> is configured as trusted, and therefore it is listed in the **antiClickJackingUri** 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 **antiClickJackingUri** 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.

8.2.8.2 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>
```

8.2.8.3 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 **applicationContext-security-filters.xml** file in a text editor. The file is located in <MashZone NextGen installation>/apache-tomcat/webapps/[presto|mashzone]/WEB-INF/classes.
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>**

8.2.8.4 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>
```

8.2.8.5 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>
```

8.2.9 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 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 is accounted for in future audits.

You can remove this personal data from your MashZone NextGen system for General Data Protection Regulation (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 Windows® Explorer, go to the directories mentioned above and delete all relevant log files. The user data is deleted from the MashZone NextGen system.

8.2.10 Create whitelists for URL calls

You can control outgoing HTTP requests when calling external URLs using URL whitelists. Use whitelists to prevent attackers from abusing HTTP requests to collect sensitive data, such as server configuration details in cloud scenarios, or to redirect users to phishing sites.

You can define URL whitelists that contain individual URLs that a user can use to access external sources. MashZone NextGen distinguishes between server-side and client-side requests based on URLs. On server-side are several data source operators that allow the data request from external URLs, for example, CSV, Excel, JSON, XML, PPM, and ARIS table. On the client-side, for example, it is the **Image** widget and the actions you can specify for multiple widgets.

Server- and client whitelists are stored in the MashZone NextGen database repository. The size of each whitelist is limited to 4000 characters. The whitelist is a JSON array of regular expressions, see the example below. URLs that match at least one regular expression have passed the check.

Example

The following JSON array represents a URL whitelist.

```
[  
  "https://www.softwareag.com.*",  
  "https://github.com/w3c/csvw.*",  
  "https://www.w3schools.com/xml.*"  
]
```

To create a URL whitelist, you can use an appropriate text editor. With the **padmin** command line tool, you can provide the created whitelist to MashZone NextGen. The padmin tool is located in the following directory:

```
<MashZone NextGen installation>/prestocli/bin/padmin.bat
```

Use the following commands to import your URL whitelist into MashZone NextGen:

- On server-side

```
padmin importBackendUrlWhitelist -f backend_whitelist.json -u Administrator -w manage -t default
```
- On client-side

```
padmin importFrontendUrlWhitelist -f frontend_whitelist.json -u Administrator -w manage -t default
```

You can also export your URL whitelist using the following commands. The server returns a JSON array of regular expressions, as shown in the above example, written to the specified file in the file system.

- On server-side

```
padmin exportBackendUrlWhitelist -f backend_whitelist.json -u Administrator -w manage -t default
```
- On client-side

```
padmin exportFrontendUrlWhitelist -f frontend_whitelist.json -u Administrator -w manage -t default
```

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 121) 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 124) 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 132) for instructions.
6. Configure the Event Service. See Event Service Configuration and Administration (page 236) for instructions.
7. If you have also installed BigMemory Max and received your BigMemory Max license, add this license to MashZone NextGen and configure MashZone NextGen to work with BigMemory Max. See Manage Licenses for MashZone NextGen and BigMemory Max (page 123) and Configure BigMemory Max 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 164) for details.

8.3.2 What is Installed with MashZone NextGen

MashZone NextGen initially installs the applications **MashZone NextGen Server** and **MashZone NextGen Event Service**.

MashZone NextGen also installs the following additional software:

- Apache's Tomcat Servlet Container, version 8.5.30
- Derby Database, version 10.13.1.1

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. For details, see the chapter Move the MashZone NextGen repository to a robust database solution (page 124).

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 Max

To use MashZone NextGen a license is required.

If you are using BigMemory Max features that require this, you also need to make your BigMemory Max license available to the MashZone NextGen Server and/or the Integrated MashZone Server. See BigMemory Max for Caching, Connections and MashZone NextGen Analytics (page 168) 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 Max 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 Max 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 121) 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 120) 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 126) for instructions.
- **MySQL:** see Move the MashZone NextGen repository to MySQL (page 128) for instructions.
- **Oracle:** see Move the MashZone NextGen repository to Oracle (page 129) for instructions.
- **PostGres:** see Move the MashZone NextGen repository to PostGres (page 131) 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, ensure that at least one user in your LDAP Directory has administrator privileges for MashZone NextGen before you move the MashZone NextGen Repository. For details, see the chapter Start MashZone NextGen with an initial administrator user (page 138).

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 120) 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. 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
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 NextGenMashZone 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. Open the MashZoneNG-install/apache-tomcat/conf/context.xml configuration file in the text editor of your choice.
 7. 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 the official Tomcat Datasource Properties 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 271).

8. 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 125) for suggestions.
9. 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 125) 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. For details, see the chapter Start MashZone NextGen with an initial administrator user (page 138).
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 the official MySQL 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 120) 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. 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.
5. Copy the MySQL JDBC driver jar file to MashZoneNG-install/apache-tomcat/lib.
6. Open the MashZoneNG-install/apache-tomcat/conf/context.xml configuration file in the text editor of your choice.
7. 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** including 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 the official Tomcat Datasource Properties 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 271).

8. 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 125) for suggestions.

Note

Depending on your local MySQL settings, MashZone NextGen may display an error message, such as "Error: Packet for query is too large (xxx > yyy)". In this case, a file (for example, the JDBC driver library) was uploaded to MashZone NextGen that exceeds the size of the **max_allowed_packet** parameter in the MySQL configuration. After setting an appropriate value in the **mysql.cnf** file, this error should no longer occur. For details, please refer the corresponding MySQL documentation.

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. For details, see the chapter Start MashZone NextGen with an initial administrator user (page 138).

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 the official Oracle documentation.

Make sure this database is supported by MashZone NextGen. See Additional MashZone NextGen System and Software Requirements (page 120) 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. 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.
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. Open the MashZoneNG-install/apache-tomcat/conf/context.xml configuration file in the text editor of your choice.
7. 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 the official Tomcat Datasource Properties 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 271).
8. 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 125) for suggestions.
9. 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 125) 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. For details, see the chapter Start MashZone NextGen with an initial administrator user (page 138).

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 the official PostgreSQL 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 120) 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. 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.
5. 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>
```
6. 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`.
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="org.postgresql.Driver" url="jdbc:postgresql://host-
name:port/databasename" />
```

- b. The JTA managed property **must** be **false**.
- c. If needed, update optional properties. See the official Tomcat Datasource Properties 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 271).

9. 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 125) for suggestions.

10. 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 125) 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 121) 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 218) 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 218) 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 133).

Authentication mechanisms. See [Defining the Authentication Scheme](#) (page 134).

Authorization mechanisms. See [Defining the Authorization Scheme](#) (page 135).

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 136).

5. Ensure that there is at least one user with administrative permissions in MashZone NextGen, otherwise the system will be locked. For details, see the chapter [Start MashZone NextGen with an initial administrator user](#) (page 138).
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.

8.3.7.1 Defining LDAP Connection Configuration

Procedure

1. Open the MashZone NextGen Admin Console (page 106).

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 134), Defining the Authorization Scheme (page 135) and Enabling MashZone NextGen Application Queries for All LDAP Users or Groups for Permissions (page 136)).
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 106).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.
3. 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.

Your existing LDAP users and groups can 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.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.
3. 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.

4. 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 **jwalker**, MashZone NextGen would search all entries in **ou=groups** where **uniquemember=jwalker**. 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 **jwalker** 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 **jwalker**.

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 106).
2. Expand **MashZone NextGen Repositories** and click **User Repository - LDAP**.
3. 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.

4. **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

5. 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.

6. 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.

8.3.7.5 Start MashZone NextGen with an initial administrator user

You can specify an initial user with administrator permissions to grant access to MashZone NextGen. This user can manage all role assignments of users and user groups in the MashZone NextGen Admin Console. In this way you can ensure that there is at least one user with administrative permissions in MashZone NextGen.

Procedure

1. Stop the MashZone NextGen server.
2. Open the **wrapper.conf** file in an appropriate text editor.
3. Specify the user ID of the initial administrator as follows.
`wrapper.java.additional.40=-Dusermanagement.initialadminaccount=myadminuser`
<myadminuser> must be a valid user in LDAP. The user receives administrator permissions in MashZone NextGen.
4. Save your settings.
5. Restart the MashZone NextGen server.

MashZone NextGen is started with the specified user with administrator permissions.

Tip

You can add LDAP users to the MashZone NextGen user repository and assign user roles to them. For details on how to add users and groups to MashZone NextGen, and to assign the relevant roles, see chapter Use the default MashZone NextGen user repository (page 138).

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 110) 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 139)

- Manage User Groups (page 140)

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 139) and Edit, assign roles and other user management tasks (page 139) for instructions.

If you have configured MashZone NextGen to use your LDAP Directory as the User Repository (page 132), you manage users and groups in LDAP and assign permissions in MashZone NextGen. For details, see Edit, assign roles, and other user management tasks (page 139) and Manage user groups (page 140).

8.3.8.1.1 Create users

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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.

See also Edit, assign roles and other user management tasks (page 139)

8.3.8.1.2 Edit, assign roles and other user management tasks

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 add users to user groups or remove users from user groups.
 - a. Enter part of a group name, if needed, and click **Search** to see a list of user groups and the user groups currently assigned to this user.
 - b. Click a group from the left column to assign a user group to this user. Click a group from the right pane to remove a user group.
 - c. Click **Save changes** to update groups for this user.
8. Click  **Edit user roles** to grant or manage permissions for a user.
 - a. Click **Search roles** to list all available roles. Or enter a role name and click **Search roles** to search a specific role.
 - b. Click a role name in the left column to assign the role to the user. Click a role in the **Assigned Roles** box to remove a role.
 - c. Click **Save changes** to update the roles assigned to the user.
9. Click **Delete user** to delete a user and confirm. Note that there must be at least one active user with **Administrator** role assigned. Therefore it is not possible to deactivate or delete the last user with **Administrator** role assigned.

8.3.8.2 Manage user groups

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 can have one or more roles assigned.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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. Click  **Manage roles** to grant or manage permissions for this user group.

- a. Click **Search roles** to list all available roles. Or enter a part of a role name and click **Search roles** to search a specific role.
 - b. Click a role name in the left column to assign the role to the user. Click a role in the **Assigned Roles** box to remove a role.
 - c. Click **Save changes** to update the roles assigned to the user.
5. To delete a user group, click **Delete** and confirm this.

8.3.8.3 Which user roles exist?

The following roles exist in MashZone NextGen.

The different roles bundle various permissions.

- Admin Console User: has read-only access to the Admin Console.
- Administrator: has all available permissions in MashZone NextGen.
- Dashboard Designer: has all permissions to create and edit dashboards.
- Feed Designer: has all permissions to create and edit data feeds.

8.3.9 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.

8.3.10 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 142) 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 145) 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.

8.3.10.1 MashZone NextGen plug-in

8.3.10.1.1 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
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

8.3.10.1.2 Instance Configuration

The instance configuration supports four different configuration types:

- General properties contained files "prestodb.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 TLS configuration (selfsigned, anyhosts) ▪ Parameters to control DAO/DB code 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/	<ul style="list-style-type: none"> ▪ Java home and Java command locations ▪ Java Main class to be used ▪ Java classpath ▪ Wrapper log settings

Configuration type	Configuration instance	Location	Covered aspects
		custom_wrapper.conf	<ul style="list-style-type: none"> 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 150).

8.3.10.1.3 Instance Logs

The **Logs** tab provides the MashZone NextGen's log files located in the **<TOMCAT_HOME>/logs** directory.

8.3.10.1.4 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.

8.3.10.2 MashZone NextGen Event Service plug-in

8.3.10.2.1 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.

8.3.10.2.2 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 in instance	Location	Description
General Properties	Event Service Core Properties	<EventService_ROOT>/conf/rtbs.properties	EDA Configuration <ul style="list-style-type: none"> ▪ EventTypeStore location (com.softwareag.event.routing.eventtypestore.location)
General Properties	HA-Store	<EventService_ROOT>/ha-store/conf/server.properties	Persistence settings for high availability <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)

Configuration type	Configuration instance	Location	Description
			<ul style="list-style-type: none"> 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 making 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>
General Properties	Zookeeper		Apache Zookeeper configuration. It is not recommended making 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

Configuration type	Configuration instance	Location	Description
			<ul style="list-style-type: none"> 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 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.
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 NextGen MashZone 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 114)).

<EventService_ROOT> usually is **<INSTALL_ROOT>/MashZoneNG/rtbs**

8.3.10.2.3 Instance Logs

The **Logs** tab provides the MashZone NextGen Event Service log files located in the **<EventService_ROOT>/log** directory.

8.3.10.3 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 TLS mode.
If you want to run the TLS mode, you must also configure the Keystores and Truststores in Command Central (page 148).
7. Save your settings.
8. Restart the MashZone NextGen server.

Your settings are applied.

8.3.10.4 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 TLS mode in the ports configuration (page 147).

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.

8.3.10.5 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 TLS mode in the ports configuration (page 147).

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.

8.3.10.6 Configure TLS certificates

You can configure TLS certificates using the Command Central.

Procedure

1. Enable in TLS mode in the port configuration (page 147).
2. Configure a keystore (page 148).
3. Configure a truststore (page 149).

Your settings are applied.

8.3.10.7 Manage licenses

You can manage the licenses of MashZone NextGen and the optional internally used product Terracotta and BigMemory Max 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.

8.3.11 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

8.4 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 Max Server(s). 	<p>Configuration When MashZone NextGen Uses Only Heap Memory (page 151)</p>
<ul style="list-style-type: none"> ▪ You have installed BigMemory Max Server(s), or ▪ Your computer has more than 4G of RAM memory. 	<p>Configuration When MashZone NextGen Uses Heap and Off-Heap Memory (page 152)</p>

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 Max 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 121) 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 Max Servers.

You must have installed a copy of your BigMemory Max license in MashZone NextGen to use off-heap memory. See Manage Licenses for MashZone NextGen and BigMemory Max (page 123) for instructions.

With combination heap and off-heap memory, as this figure shows, BigMemory Max 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 Max 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 Max, 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 121) for instructions.

Memory options used with BigMemory Max	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 Max 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 Max is deployed.

Memory options used with BigMemory Max	Description
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 Configure feed processing time zone

MashZone NextGen feed processing represents date values in an internal format that refers to the time zone. All text to date or date to text conversions also use the configured time zone, unless an explicit time zone is specified in values or formatting functions.

The time zone is configured in the **mashzone.properties** file located in the following directory.

<MashZone NextGen installation>/apache-tomcat/webapps/mashzone/WEB-INF/

Open the **mashzone.properties** file and specify the **mashzone.feedprocessing.timezone** property.

You can only use time zones that are compatible with the Java **TimeZone.getTimeZone** method. For more details, see the corresponding Java documentation.

Examples

- UTC
- Europe/Berlin
- Asia/Kolkata
- GMT-5
- GMT+4:30

The default value is UTC.

This setting introduces unified time zone handling throughout feed processing. Note that this setting applies to all feed processing operators created from MashZone NextGen 10.4. Operators in feeds and dashboards that were migrated or imported from MashZone NextGen 10.3 and earlier, will run in MashZone NextGen 10.3 compatibility mode with respect to time zones. These feeds and dashboards use time zones as in MashZone NextGen 10.3 and earlier, that is, the **Runtime Info** operator creates a shifted date, and some RAQL functions use the server's local time zone.

8.4.3 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 155) 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 155) for details.

8.4.3.1 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 126)
 - [Move the MashZone NextGen repository to MySQL](#) (page 128)
 - [Move the MashZone NextGen repository to Oracle](#) (page 129)
 - [Move the MashZone NextGen repository to PostGRES](#) (page 131)

8.4.3.2 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 272) 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.jsfile`. 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.

8.4.3.3 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 uses 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.

8.4.3.4 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 121) for instructions.

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 121) 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 106).
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 121) 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 **white list** of addresses that do not require proxy server access. See Define a Proxy Server White list for MashZone NextGen (page 159) for more information.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 white list of addresses that should not use the proxy server. See Define a Proxy Server White list for MashZone NextGen (page 159) 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 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 NextGen-installation>/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/.*" 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 White list for MashZone NextGen

If you have configured a proxy server for the MashZone NextGen Server, you can define a white list of domains, hosts or IP addresses that do **not** require access through the proxy server.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 White list (page 160).

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 White list (page 160).

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 White list (page 160).
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 White list

All of the white list 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 160)
- Specifying Domains (page 161)
- Specifying Host Names (page 161)

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 white list 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 white list 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 160) for more information.

8.4.8 Configure MashZone NextGen for TLS and Digital Certificates

MashZone NextGen expects HTTP as the default transport protocol from clients to the MashZone NextGen Server.

MashZone NextGen supports HTTPS and TLS for connections from clients or connections to many types of information sources.

The certificate store, certificates and configuration needed to support TLS 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 TLS to MashZone NextGen (page 163).	✓			✓			
Mutual to MashZone NextGen See Configure Mutual TLS Between Users and MashZone NextGen (page 163).	✓	✓		✓		✓	
One-Way TLS to Mashable Information Sources (page 163).		✓	can be in either				

See also The Certificate Store and Certificates (page 162) 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 TLS 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 TLS. The key certificate pair identifies the MashZone NextGen Server to information sources for mutual TLS.

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 TLS, or identify information sources, for one-way or mutual TLS.

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 TLS 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 162) for more information.
 - See The Certificate Store and Certificates (page 162) 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 164) for instructions.

8.4.8.6 One-Way TLS to MashZone NextGen

This requires:

1. A key store and a key certificate pair for MashZone NextGen. See The Certificate Store and Certificates (page 162) 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 164) for instructions.

8.4.8.7 One-Way TLS to Information Sources

This requires:

- A trust store for MashZone NextGen. See The Certificate Store and Certificates (page 162) 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 164) for instructions.

- Java. See Update TLS Configuration for Java (page 165) for instructions.
- Self-signed certificates, if any, for data source using one-way TLS. You must add these certificates to the trust store before the data source can be registered. See Trusted Peer Certificates (page 162) for more information.

8.4.8.8 Configure HTTPS and Certificate Stores in the Application Server

Configuration for TLS for MashZone NextGen can be defined in the application server that hosts the MashZone NextGen Server. These instructions discuss the basic steps for configuring TLS in Tomcat. See the official Tomcat documentation 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 162) for instructions.
2. Configure Tomcat for secure connections from clients to the MashZone NextGen Server. We recommend using TLSv1.2 for network communication.

Warning

We recommend not using TLSv1.0 and TLSv1.1 in your production environments, as TLSv1.0 and TLSv1.1 are considered unsafe for network communication.

You can add the **sslEnabledProtocols** and **sslProtocols** parameters to the **server.xml** file to disable TLSv1.0 and TLSv1.1. The parameters allow limiting the supported protocols. For details, see the official Tomcat documentation.

- a. Edit the server.xml file for Tomcat to uncomment and configure the <Connector> element for TLS/HTTPS 1.1. For example:

```
<Connector port="8443"
protocol="org.apache.coyote.http11.Http11Nio2Protocol"
SSLEnabled="true" maxHttpHeaderSize="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 TLS, 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 156) 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 TLS 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 121) 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 166) for more information.
- Audit logging for dashboards, data feeds, aliases, and permissions. See Audit logging for dashboards, data feeds, aliases, and permissions (page 167) 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 166) 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 BigMemory Max for Caching, Connections and In-Memory Stores

By default MashZone NextGen uses local memory for caching. This uses BigMemory Max 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 Max Servers on one or more additional hosts and configure MashZone NextGen and the Integrated MashZone Server to work with them. MashZone NextGen requires BigMemory Max Servers with a BigMemory Max license to support:

- Significant, extensible amounts of memory, most commonly for very large datasets used with MashZone NextGen.

BigMemory Max 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 Max Servers.

- Access to off-heap memory

BigMemory Max 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 Max 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 Max connections as a data source.

See [Configure BigMemory Max Servers for MashZone NextGen Caching and In-Memory Stores](#) (page 170) for instructions on configuring MashZone NextGen to work with BigMemory Max 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 Max Stores for MashZone NextGen](#) (page 171)
- [Manage Dynamic BigMemory Max Stores for MashZone](#) (page 176).

See [Caching for the MashZone NextGen Server](#) (page 169) for an overview of MashZone NextGen caching. For caching configuration, see [Distributed Caching for MashZone NextGen Clusters](#) (page 169).

8.4.10.1 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.

8.4.10.1.1 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.

8.4.10.1.2 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.

8.4.10.1.3 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 Max Servers.

You use BigMemory Max Servers to handle distributed caching for MashZone NextGen.

With BigMemory Max 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 Max 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 Max Servers for MashZone NextGen Caching](#) (page 170) for set up instructions.

8.4.10.1.4 Configure BigMemory Max Servers for MashZone NextGen Caching

You can configure MashZone NextGen to work with one or an array of BigMemory Max Servers to provide additional memory, provide reliability and support specific other features. See [BigMemory Max for Caching, Connections and In-Memory Stores](#) (page 168) for more information on features that require BigMemory Max servers.

Procedure

1. Copy your license for BigMemory Max 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 Max server(s) you installed. For example:

```
<terracottaConfig url="tcHost1:9510" />
```

There are several ways to identify one or more BigMemory Max servers for MashZone NextGen. See BigMemory Max documentation (<http://terracotta.org/documentation/>) for more information.

- 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 Max hosts. This combined memory is managed by BigMemory Max.

For more information on the `<terracotta>` element, see **Distributed Configuration** topics in BigMemory Max documentation (<http://terracotta.org/documentation/>).

- 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.

- Start BigMemory Max Server(s).
- If needed, adjust memory configuration for the local MashZone NextGen host. See Memory Configuration for the MashZone NextGen Server (page 150) for instructions.
- Restart MashZone NextGen. See Start and Stop the MashZone NextGen Server (page 121) for instructions.

8.4.10.2 Working with BigMemory Max Stores for RAQL

Declaring In-Memory Stores for MashZone NextGen allows MashZone NextGen to connect to stores in BigMemory Max 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 Max 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.

8.4.10.2.1 Declared In-Memory Stores

Declared stores are defined in BigMemory Max 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).

8.4.10.2.1.1 Declare a new In-Memory Store

Define configuration for one or more In-Memory Stores in a cache configuration file for BigMemory Max (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 Max 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 Max Search API or an implementation class of **AttributeExtractor**. See BigMemory Max 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 Max.

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 218).

5. Restart the MashZone NextGen Server. For instructions, see [Start and Stop the MashZone NextGen Server](#) (page 121).
6. Open the MashZone NextGen Admin Console (page 106).
7. Click **Server** to expand this section of the **Administration** menu.
8. Click **BigMemory Max**.
9. Open the **BigMemory Max Cache** tap.

The **BigMemory Max 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 BigMemory Max 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 196) for more information.

8.4.10.2.1.2 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 Max hosts when you add servers or deploy MashZone NextGen in staging or production environments.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the Administration menu.
3. Click **BigMemory Max**.
4. Open the **BigMemory Max Cache** tap.
5. The **BigMemory Max Cache** tab lists any In-Memory Store configuration files that have already been upload.
6. Select the existing BigMemory Max 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.

8.4.10.2.1.3 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 106).
2. Click **Server** to expand this section of the Administration menu.

3. Click **BigMemory Max**.
4. Open the **BigMemory Max 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.

8.4.10.2.1.4 Define permissions for declared In-Memory Stores

You can define permissions for declared In-Memory Stores.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **BigMemory Max**.
4. Open the **BigMemory Max Cache** tap.

The tab lists all available configuration files for declared In-Memory Store that have been upload.
5. Click the  **Edit permission** icon to configure the permission of an existing EhCache configuration file.
6. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
7. Drag a user or user group from the **Search results** field and drop it into the **Principals with permissions** field.

By default, the owner of the EhCache configuration file is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.
8. Enable or disable the **View** or **Usage** privileges of a user or user group.
9. Click **OK**.

Your settings are applied.

8.4.10.3 Dynamic In-Memory Stores

Dynamic stores are created programmatically on the fly when:

- An external system creates a store dynamically in BigMemory Max. 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 Max be installed as a server or server array. MashZone NextGen must also have access to the BigMemory Max 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 Max 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 Max API. See [BigMemory Max 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.

8.4.10.3.1 Manage Dynamic BigMemory Max Stores for MashZone NextGen Analytics

You must define connections and identify configuration information for BigMemory Max 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 Max Server.

You can also define connections to external in-memory stores that are not created dynamically. See [Declare BigMemory Max 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 176) and [Delete External Dynamic In-Memory Store Connections](#) (page 177).

8.4.10.3.2 Add an External Dynamic In-Memory Store Connection

Procedure

1. Verify that the Terracotta Management Console (TMC) that manages the BigMemory Max 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 218).

4. Restart the MashZone NextGen Server. For instructions, see Start and Stop the MashZone NextGen Server (page 121).
5. Open the MashZone NextGen Admin Console (page 106).
6. Click **Server** to expand this section of the Administration menu.
7. Click **BigMemory Max**.
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 BigMemory Max 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 BigMemory Max.
Cache Manager names are a best practice for dynamic external stores. If the external system does not assign a cache manager name, BigMemory Max uses a default name which can lead to name collisions and errors.
14. If the TMC requires TLS for connections, change the Security type to TLS.
15. You can enter an user name and a password optionally.
16. Click **Add this external cache** to save this connection.

8.4.10.3.3 Delete External Dynamic In-Memory Store Connections

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the Administration menu.
3. Click **BigMemory Max**.
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.

8.4.10.3.4 Define permissions for external dynamic In-Memory Store Connections

You can define permissions for external dynamic In-Memory Store connections.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **BigMemory Max**.
4. Open the **Terracotta Management Server** tap.

This tab lists connections to any existing Terracotta Management Server.
5. Click the  **Edit permission** icon to configure the permission of an existing Terracotta Management Server.
6. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
7. Drag a user or user group from the **Search results** field and drop it into the **Principals with permissions** field.

By default, the owner of the Terracotta Management Server alias is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.
8. Enable or disable the **View** or **Usage** privileges of a user or user group.
9. Click **OK**.

Your settings are applied.

8.4.11 Manage Terracotta connections

You can register, edit, delete, and test Terracotta connections. Additionally, you can share Terracotta connections with particular users and user groups.

8.4.11.1 Register Terracotta connections

You can configure a connection to the Terracotta server. The alias is used to refer to that specific connection.

By having a connection to the Terracotta server configured, the dashboard developer can access all datasets available on that server. The datasets are provided in the corresponding **Terracotta** source operator. The alias is unique for all Terracotta connections. Different connections can point to the same Terracotta server.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).

2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta**.
4. Click **Register new Terracotta connection**.
5. Enter a name for the Terracotta connection in the **Terracotta alias** field. The connection data is saved under this alias.
6. Enter the URI to the Terracotta. The format is `terracotta://hostname:port`.
7. You can select the security type of the Terracotta connection in the **Security type** drop-down menu. If the Terracotta connection is secured, select the **Secured** security type. The default value is **None**. For details, see [Ensure a secure interaction with Terracotta \(page 181\)](#).
8. Enter the URI to the security root directory in the corresponding input box. The input box is available if **Secured** is selected in the **Security type** drop-down menu.
The directory provides the artifacts that MashZone NextGen requires to securely interact with the Terracotta server.
9. Enable the **Materialized Feed support** option if you want to manage Materialized Feeds with this Terracotta connection.
 - a. Specify the **Off-heap resource** configured on the Terracotta server that you want to use to materialize the data feed results.
 - b. Terracotta provides a fast restart capability by writing the contents of a dataset in addition to the hard disk. If you want to use this feature, specify the name of the disk resource configured on the Terracotta server.
10. Click **Add connection**.

The **Terracotta** connection is created and listed by alias name.

Click  **Test connection** for the created alias to test the Terracotta server connection.

8.4.11.2 Edit Terracotta connections

You can edit the URI of an already existing connection to the Terracotta server. The alias is not editable.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta**.
4. Click  **Edit**.
5. Enter the URI to the Terracotta. The format is `terracotta://hostname:port`.
6. You can select the security type of the Terracotta connection in the **Security type** drop-down menu. If the Terracotta connection is secured, select the **Secured** security type. The default value is **None**. For details, see [Ensure a secure interaction with Terracotta \(page 181\)](#).

7. Enter the URI to the security root directory in the corresponding input box. The input box is available if **Secured** is selected in the **Security type** drop-down menu.

The directory provides the artifacts that MashZone NextGen requires to securely interact with the Terracotta server.

8. Enable the **Materialized Feed support** option if you want to manage Materialized Feeds with this Terracotta connection.
 - a. Specify the **Off-heap resource** configured on the Terracotta server that you want to use to materialize the data feed results.
 - b. Terracotta provides a fast restart capability by writing the contents of a dataset in addition to the hard disk. If you want to use this feature, specify the name of the disk resource configured on the Terracotta server.
9. Click **Update connection**.

Your changes are applied.

Click ► **Test connection** to test the Terracotta server connection.

8.4.11.3 Test Terracotta connections

You can test the connection to the Terracotta server to check if the connection works correctly.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta**.
4. Click ► **Test connection** for an alias to test the Terracotta server connection.

If the connection is used for Materialized Feeds, the provided values for off-heap and optional disk resources are tested as well.

The server connection is tested and a test result is displayed.

8.4.11.4 Share Terracotta connections

You can share Terracotta connections with particular users and user groups so that they have access to the Terracotta dataset and manage datasets with feed results.

Regardless of the users and user groups with whom you have shared the connections, users with administrative privileges can access all Terracotta aliases.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta**.
4. Click the  **Edit URL alias permissions** icon of the Terracotta connection you want to share.

5. Enter a term in the search field and click **Search**. Clicking **Search** without any input values retrieves all users and groups.
6. Drag a 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 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 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 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 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 connection, and the **Display** privilege for the corresponding data feed.
 See also Share Materialized Feeds (page 192).
8. Click **OK**.

Your changes are applied.

8.4.11.5 Delete Terracotta connections

You can delete a Terracotta server connection.

Warning

A deleted Terracotta server connection cannot 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 106).
2. Click **Server** to expand this section of the **Administration** menu.
3. Click **Terracotta**.
4. Click  **Delete**.

The selected Terracotta connection is deleted.

8.4.11.6 Ensure secure interaction with Terracotta

MashZone NextGen interacts as a Terracotta client with the Terracotta server. The Terracotta server uses the following security features to ensure a secure connection to MashZone NextGen.

- Connection encryption: The connections between client and server are encrypted using the SSL/TLS protocol.

- Cluster authentication: The identity of clients connecting to a server in a Terracotta cluster is validated using certificate-based, file-based, or LDAP-based authentication.

The security features in Terracotta are configured using a security root directory. The security root directory is a file system location for all security-related files, such as certificates. The client as well as the server must provide a security root directory. Both must match in their security configuration, that is, if the server has been configured for LDAP in its security root directory, the client must also have a corresponding LDAP configuration file in its security root directory. Therefore, MashZone NextGen as a Terracotta client must point to a security root directory in order to securely interact with a Terracotta server. This security root directory must contain the appropriate security files in advance and these files must match the security configuration of the server. Detailed information on the various security configurations and the structure of a security root directory can be found in the Terracotta documentation.

Because different Terracotta connections configured in MashZone NextGen can point to Terracotta servers with possibly different security configurations, each Terracotta connection in MashZone NextGen must point to a separate security root directory with a matching security configuration. Good practice is to manage all required security root directories in one central location. This location must be <MashZone NextGen installation>/apache-tomcat/conf/terracotta-security.

Ensure that the security root directories used by MashZone NextGen are accessible only by the user who is permitted to run the MashZone NextGen server.

8.4.12 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 182), Edit, test or remove data sources (page 184) and Add or manage JDBC drivers (page 185) for instructions.

8.4.12.1 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 106).
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 185) 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**.

8.4.12.2 Edit, test or remove data sources

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 182) 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.

8.4.12.3 Share data sources

You can share data sources with particular users and user groups so that these have access to **JDBC data sources**.

Prerequisite

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 106).
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. 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.

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.

8.4.12.4 Add or manage JDBC drivers

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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**.

8.4.12.5 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.

8.4.12.5.1 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.

8.4.12.5.2 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 NextGenMashZone 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.

8.4.12.5.3 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.

8.4.12.5.4 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.

8.4.13 Manage geographical map resources

8.4.13.1 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 with 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 301) 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-continents-by-name
- us-states-by-name

If you add your own map style templates, we recommend that all defined regions (features) are 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. Therefore, any polygon without an **ID** attribute is also ignored. If no data is assigned, you can use the vector map to select a region and publish such a selection using the region ID.

8.4.13.2 Manage tile server configuration files

TMS (Tile Map Service) is a protocol for serving maps as tiles, which involves 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}\tileserver

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 have been included as a point of reference.

The maps are available in the **Map with markers** (page 292) 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. 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"]
}
```

8.4.14 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 189).
- Manual tuning gives you control to balance memory requirements for MashZone NextGen, but does require manual updates to several configuration files.

8.4.14.1 Tune MashZone Memory and Cache Configuration Manually

To manually update memory and cache configuration

- Update Cache Memory Settings (page 190)
- Update MashZone ThreadSize Properties (page 190).

- Then restart the MashZone NextGen Server to apply this change. See Start and Stop the MashZone NextGen Server (page 121) for instructions.

8.4.14.2 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

8.4.14.3 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

8.4.15 Manage Materialized Feeds

8.4.15.1 Configure Materialized Feeds

Using Materialized Feeds (page 97), you can materialize the results of a data feed evaluation in Terracotta.

To work with those results, you can run RAQL queries over the corresponding dataset using the Terracotta (page 317) data source operator.

Prerequisite

You have configured a Terracotta connection to use materialized data feeds.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 connection in the **Terracotta connection** drop-down menu. The menu contains all Terracotta connections supporting Materialized Feeds.

The dataset for materializing feed results will be created on the Terracotta server defined for this Terracotta connection.

7. Enter a name for the dataset used for materializing feed results in the **Dataset name** input box.

Either use the predefined name or specify a custom name. The name must begin with the prefix **MZNG**. The dataset is created if it does not already exist. 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.

8.4.15.2 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 106).

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 connection in the **Terracotta connection** drop-down menu. The menu contains all Terracotta connections supporting Materialized Feeds.

The dataset for materializing feed results will be created on the Terracotta server defined for this Terracotta connection.

5. Enter a name of the dataset used for materializing feed results in the **Dataset name** input box.

A new, empty dataset is created with the new dataset name. The old dataset is not deleted.

6. Click **Update**.

Your changes are applied.

8.4.15.3 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 106).
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.

8.4.15.4 Share Materialized Feeds

You can share Materialized Feeds with particular users and user groups so that they 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.

Users with administrative privileges can manage all Materialized Feeds.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 retrieves all users and groups.
5. Drag a 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 not editable and cannot be removed from the list.
6. Enable or disable the **Display** or **Usage** privileges of a user or user group.
7. 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 connection that is used. For details, see Which privileges are required for Materialized Feeds? (page 100)

8.4.15.5 Materialize data feed results in Terracotta

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 dataset.

You cannot run this process 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 you start a new run, 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.

8.4.15.5.1 Use the Admin Console

You can materialize data feed results in Terracotta using the MashZone NextGen Admin Console.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 connection.

After the feed calculation and the subsequent materialization of results are completed, you can access the feed results using the Terracotta (page 317) source operator.

Tip

Additional information, such as the duration or potential execution errors, is available in the MashZone NextGen log file.

8.4.15.5.2 Use the server administration tool

You can materialize data feed results in Terracotta 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 you do not specify input parameters, default values as defined in the data feed are used.

-u username: Is the MashZone NextGen username with which to log in. This account must have MashZone NextGen administrator permissions.

-w password: Is the MashZone NextGen password with which to log in.

The data feed results are calculated and materialized using the configured Terracotta 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 (page 317) source operator.

Tip

Additional information, such as duration or potential execution errors, is available in the MashZone NextGen log file.

8.4.15.5.3 Use server web services

You can materialize data feed results in Terracotta 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.

For **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 the above authentication cookie, you can use the public logout service to finish your session.

For **Authentication using basic credentials**, see below.

The data feed results are calculated and materialized using the configured Terracotta connection.

After the feed calculation and the subsequent materialization of results is completed, you can access the feed results using the Terracotta (page 317) source operator.

Tip

Additional information, such as duration or potential execution errors, is available in the MashZone NextGen log file.

AUTHENTICATION USING BASIC CREDENTIALS

For authentication using basic credentials, you must use the REST API method to log in 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 log out 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
```

8.4.16 Disable automatic masking in CSV files

You can disable the automatic masking when you export data as a CSV file.

By default, all values beginning with =, +, -, or @ are automatically masked by a preceding single quotation mark before exported as CSV file. For details on exporting data as a CSV file, see the chapter Save widget data as a CSV file (page 10).

Procedure

1. Open the **mashzone.properties** file in the following directory.
<MashZone NextGen installation>/apache-tomcat/webapps/mashzone/WEB-INF/
2. Set the **allow.secure.csv.export parameter** to false as follows.
allow.secure.csv.export=false

The automatic masking when exporting CSV files is disabled.

8.5 MashZone NextGen Server Administration

8.5.1 Manage files for MashZone NextGen dashboards and data feeds

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 196)
- Find MashZone NextGen Files (page 197)
- Update or delete MashZone NextGen files (page 197)
- **Share MashZone NextGen resource files (page 198)**

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 106).
2. **Expand the Platform Features section and click File Resources.**
3. Click **Upload New File**.
4. Click **File**, find and select the file you want to upload and click **Open**.
The location and file name are filled 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 default file name is **/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 a slash `/` 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.

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

You can search for specific resource files.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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**.

The search result is displayed. The 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 106).
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 197) for techniques.
4. To upload an updated file:
 - a. Click  **Edit** on the line for that file.
 - b. Click **File** 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.

8.5.1.4 Share MashZone NextGen resource files

You can share resource files with particular users and user groups so that these have access to the files. MashZone NextGen users can use these files as external resources in their dashboards and data feeds.

Regardless of the share, users with administration privilege can access all resource files. By default, all other users only have access to images (such as PNG, JPG, or GIF).

Prerequisite

You have administration privileges.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Expand the **Platform Features** section and click **File Resources**.
3. Click  **Edit file permission** to configure the file permission for an existing file resource.
4. Enter a term in the search field and click **Search**. Clicking on **Search** without any input values fetches all users and groups.
5. Drag a user or user group from the **Search results** field and drop it into the **Principals with permissions** field.

By default, the owner of the resource file is already present in the **Principals with permissions** list. This owner is non editable and cannot be removed from the list.

6. Enable or disable the **Usage** privileges of a user or user group.

A user or user group with **Usage** privilege has access to the corresponding resource file and can use it in data feeds or dashboards.

7. Click **OK**.

Your settings are applied.

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 106).
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 106).
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 106).
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.

Regardless of the share, users with administration privilege can access all resource directories.

Prerequisite

You have administration privileges.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **File alias** -> **File 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. 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.

6. Click **OK**.

Your settings are applied.

8.5.3 Manage URL aliases

You can manage your URL aliases in the **Admin console**.

It is always recommended using a URL alias to shorten the link used in e.g. 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 106).
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 106).
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 106).
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.

Prerequisites

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 106).
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. 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.

6. 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.

7. 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 202), shown below, and optionally Deploying MashZone NextGen Artifacts and Other Metadata (page 202).

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 124) for instructions.

You can leave 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 Max. If additional memory or reliability is required, you can also deploy BigMemory Max 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 212). To deploy MashZone NextGen servers in clusters, see Clustering MashZone NextGen Servers (page 214) for requirements and links.

8.5.4.2 Deploying artifacts and other metadata using the command line

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 MashZone NextGen migration utility instead.

For details on using the migration utility, please refer the MashZone NextGen Migration Guide.

- 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, groups, and role assignments (page 204)

Export dashboards (page 205)

Export data feeds (page 206)

Export aliases (page 207)

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 218) 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 NextGen MashZone NextGen Server.

See Manage data sources and drivers (page 182) 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 (page 210)

Import data feeds (page 211)

Import users, user metadata and groups (page 208)

Import aliases (page 211)

Table 1. Known Export/Import Limitations

	Exported	Not Exported
--	----------	--------------

<p>Related Metatdata/ User Metadata</p>	<ul style="list-style-type: none"> ▪ 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 dashboards and data feeds. ▪ Datasources must be added to the target MashZone NextGen Server before you import any artifacts that use them or the import will fail.
<p>MashZone NextGen Server Configuration</p>		<p>Configuration for the MashZone NextGen Server.</p>

8.5.4.2.1 Export users, groups, and role assignments

You can export all users, groups, and role 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

1. 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 121) for instructions.
2. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.
3. Enter this command:

```
padmin exportUsersAndGroups -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/**.
- 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.

8.5.4.2.2 Export role assignments for users and groups

You can export all role assignments for users and groups 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

1. 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 121) for instructions.
2. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.
3. Enter this command:

```
padmin exportRoleAssignments -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/**.

-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.

8.5.4.2.3 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 identify [-f output-file] [-l prestoURL] -u username -w password [-v] [-o]
```

-i identify: Mandatory dashboard identifier. It can be "id=", "name=" or "all", enclosed in quotes.

Examples

-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/**.

-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.

8.5.4.2.4 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 identify [-f output-file] [-l prestoURL] -u username -w password [-v] [-o]
```

Examples

-i identify: 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/**.

-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.

8.5.4.2.5 Export aliases

You can export your MashZone NextGen aliases.

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

Procedure

1. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.

2. Enter this command:

```
padmin exportAlias -i identify [-f output-file] [-l prestoURL] -u username -w password [-o] [-j]
```

-i identifier: mandatory alias identifier. Supported identifiers are: type, name or all
Supported types: PPM, FILE, DATABASE, EVENT, TERRACOTTA, MATERIALIZED_FEED, BIGMEMORY_CONNECTION

Examples

- i "type=PPM": Aliases with type **PPM** are exported.
 - i "name=UMG_EN": Aliases with name **UMG_EN** are exported.
 - i "all": Export of all aliases 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:
Export creates a new file alias-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/.
- 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.
- j** is an optional flag to turn on inclusion jdbc driver JAR files in export.
- 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 aliases to MashZone NextGen.

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

There is also an information about the alias creator stored in the ZIP file.

8.5.4.2.6 Import users, groups, and role assignments

You must have a users export file to import. For details, see the chapter Export users, groups, and role assignments (page 204).

This command is available only in MashZone NextGen 3.2 or later.

A user export file contains MashZone NextGen User Attributes. It can also contain users, user groups, and role 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 121) for instructions.
2. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.
3. Enter this command:

```
padmin importUsersAndGroups -f input-file [-l prestoURL] -u username -w password [-c] [-o] [-v]
```

-f input-file: is the path and name of the export file to import data from.

-l 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/**.

-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.

8.5.4.2.7 Import role assignments for users and groups

You must have a users export file to import. For details, see the chapter Export role assignments for users and groups (page 205).

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 121) for instructions.

2. Open a command window and move to the **<MashZone NextGen installation>/prestocli/bin** folder.

3. Enter this command:

```
padmin importRoleAssignments -f input-file [-l prestoURL] -u username -w password [-c] [-o] [-v]
```

-f input-file: is the path and name of the export file to import data from.

-l 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/**.

-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.

8.5.4.2.8 Import dashboards

You can import dashboards in MashZone NextGen.

The dashboards are saved in a ZIP file 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 <MashZone NextGen installation>/prestocli/bin folder.
2. Enter the following command:

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

-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: 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: Optional. Enables the verbose logging.

-c: Optional. Checks whether the selected file can be imported successfully.

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.

8.5.4.2.9 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.

-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.

8.5.4.2.10 Import aliases

You can import aliases to MashZone NextGen.

The aliases are saved in a ZIP file that contains among other things the alias definitions and permissions.

Procedure

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

```
padmin importAlias [-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 permissions saved in the import ZIP file.
 - o**: is optional. Allows to overwrite an existing alias.
 - 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 aliases in MashZone NextGen.

8.5.4.2.11 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 214) for information.

8.5.4.3 Deploying artifacts using the Admin Console

You can export and import dashboards, data feeds, and aliases using the MashZone NextGen Admin Console.

8.5.4.3.1 Export dashboards, data feeds, and aliases

You can export dashboards, data feeds, and aliases using the MashZone NextGen Admin Console, for example, to share your dashboards with other users.

You can export all existing artifacts of a type at once or individual artifacts. The artifacts are exported including the corresponding permissions. The exported artifacts are downloaded and stored in ZIP files in the client system.

Prerequisites

You have administration privileges.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Import/Export** -> **Import/Export** in the **Administration** menu.
3. Select the artifact type you want to export from the **Type** drop-down menu.
4. To additionally export the relevant JDBC drivers for the aliases, enable the **Include JDBC drivers** option. The option is only available for aliases.
5. Click **Export all** to export all available artifacts of the selected type.
All artifacts are stored in one ZIP file.
6. To export a specific artifact, you can specify the corresponding name, or for dashboards and data feeds the GUID. For aliases, you can specify the alias type instead of the GUID, for example, aliases of type EVENT or PPM.
 - a. Enable the **Name** option if you want to identify the artifact by name.
Note that there may be multiple dashboards or data feeds with the same name. A name of an artifact is not unique in MashZone NextGen. All dashboards or data feeds with the same name are exported.
 - b. Enable the **GUID** option if you want to identify the artifact by GUID.
A GUID is unique in MashZone NextGen.
 - c. Enable the **Type** option if you want to identify the alias by type.
All aliases of the same type, such as EVENT or PPM, are exported.
 - d. Enter the name, GUID, or type in the **Identifier** input box.
An automatic completion function helps you to find the desired artifact and shows the relevant artifacts.
7. Click **Export**.

The selected artifacts are exported.

8.5.4.3.2 Import dashboards, data feeds, and aliases

You can import dashboards, data feeds, and aliases using the MashZone NextGen Admin Console.

The artifacts are stored in a ZIP file containing the artifact definitions and permissions, etc. If you import an artifact including permissions, the creator of the artifact can view and edit the artifact. The importer of an artifact automatically becomes the creator of the artifact if the artifact is imported without permissions.

Prerequisites

You have administration privileges.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Import/Export** -> **Import/Export** in the **Administration** menu.
3. Select the artifact type you want to import into MashZone NextGen.
4. Enable the **Overwrite if existing** option to overwrite the same already existing artifacts. By default, the import is rejected if the artifacts already exist in MashZone NextGen.
5. Enable the **Import permissions** option to also import the permissions of the selected artifacts.
6. Click **Select a ZIP file** to select the file containing the artifacts you want to import.
7. Click **Test run** to start a test run, to check for possible conflicts with different MashZone NextGen editions.
8. Click **Import** to import the selected artifacts.

The selected artifacts are imported.

When the import is completed successfully, you can use the imported artifacts in MashZone NextGen.

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 214) or Adding New Members to an Existing Clusters (page 216) 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 218): 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 217): 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 Max Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for instructions on how to configure BigMemory Max, 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 216) 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 \(page 202\)](#) 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 \(page 217\)](#) 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 \(page 219\)](#) 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 \(page 219\)](#) 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.

7. 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 217) with cluster members, typically when you are creating new environments. Or you can Share an Existing MashZone NextGen Repository (page 217) within a cluster.

8.5.6.1 Create and Share a New MashZone NextGen Repository

To create a new shared repository

1. Create a new MashZone NextGen Repository in the appropriate database for your environment. See the table below.
2. Copy the JAR file for the JDBC driver for your database to the **<MashZone NextGen installation>/apache-tomcat/lib** folder on all cluster nodes.
3. 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 124) for instructions.
4. Enable distributed caching for artifacts. See Configure BigMemory Max Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for more information and instructions.
5. If the MashZone NextGen Repository is hosted in Microsoft SQL Server, MySQL or Oracle, change the repository JAR in the MashZone NextGen Server.
6. Restart each MashZone NextGen Server in the cluster.

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 218) 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 Max 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 219) 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 219) and MashZone NextGen Extensions (page 221) 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 220) 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/mashzone/WEB-INF/classes
ehcache.xml	Configuration information for MashZone NextGen caches. This also contains	

File	Description and Configuration	Default Location
	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 BigMemory Max license file	The license file for BigMemory Max, 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 Max license in an external folder shared across the cluster. See Configure BigMemory Max Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for required configuration steps to enable a shared license.	
userRespositoryLdap.properties	Connection information for your LDAP Directory. See Integrate Your LDAP Directory with MashZone NextGen (page 132) 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
------	-------------------------------

<p>applicationContext-commonServices.xml</p>	<p>You edit configuration in this file if you choose to use distributed response caching for MashZone NextGen. See Configure BigMemory Max Servers for MashZone NextGen Caching and In-Memory Stores (page 170) for more information.</p> <p>You may need to update this configuration, as needed, to add additional distributed cache nodes to tune performance.</p>
<p>applicationContext-security.xml</p>	<p>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 110) for more information.</p>
<p>log4j.properties</p>	<p>This file is updated automatically when you change logging configuration in the Admin Console. See Configure Logging for the MashZone NextGen Server (page 166) for details.</p> <p>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.</p>
<p>userRepositoryApplicationContext.xml</p>	<p>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 132) for details.</p>
<p>userRepositoryApplicationContext-ldap.xml</p>	

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
<p>Custom filter classes for single sign-on authentication. See Implementing a Custom SSO Filter (page 112) for details.</p>	<p><MashZone NextGen installation>/apache-tomcat/ mashzone/WEB-INF/classes or</p>
<p>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.</p>	<p><MashZone NextGen installation>/apache-tomcat/ mashzone/WEB-INF/lib (for JARs)</p>

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 124).
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 dashboards 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 sub 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 configuring 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 initializations 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.

Edit the **`mashzone.feedprocessing.timezone`** property in the **`mashzone.properties`** file and specify a common timezone in all nodes of the cluster.

For more details, see the chapter **Configure feed processing time zone** (page 153).

Example

```
mashzone.feedprocessing.timezone=UTC
```

Use UTC, or replace it by the desired time zone ID. You can only use time zones that are compatible with the Java **`TimeZone.getTimeZone`** method.

Note that date values stored in Materialized Feeds will differ between nodes if time zones are configured differently within the cluster. This is because their representation is relative to the configured time zone.

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.

8.5.9 Customize application and dashboard styles

You can customize the styles of the MashZone NextGen application, dashboards, and widgets. For example, you can replace the brand logo and the text displayed on the MashZone NextGen welcome page. You can edit the look and feel of the dashboards and widgets, for example, colors schemes, fonts, or background colors. You can use the available styles in MashZone NextGen to create your own style templates.

MashZone NextGen provides two modes to customize the application and dashboard styles. For details, see the chapter Switch the style editing mode (page 231).

8.5.9.1 Customize dashboard and widget styles

You can create and add your own styles to customize the look and feel of your dashboards and widgets.

Tip

You can use the **default** style as template for your own styles. MashZone NextGen styles are stored in the database and can be downloaded and saved as LESS style files. (page 228)

In file-based mode (page 231), you can edit the LESS files stored in the following folder.

```
<MashZone NextGen installation\apache-  
tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\dashboard
```

Prerequisite

You have administration privileges.

Procedure

1. Open a LESS style file in an appropriate text editor.
2. Specify your style settings.
3. Save your settings.
4. In database mode, upload the modified style file (page 229). You can replace an available style or add a new style. The styles are uploaded if they can be compiled by the MashZone NextGen server. In database mode, the styles are automatically compiled by default.
5. In file-base mode, perform the following steps.
 - a. If you want to modify an available style, replace the corresponding LESS file with the modified file.

- b. If you want to add a new style, create a folder as described in the chapter How are the style files structured? (page 233), and store the modified LESS file in the new folder with a new file name.
- c. Compile the modified style. (page 230)

The new style is now available in the MashZone NextGen.

In the dashboard editor, you can select the uploaded styles to change the look and feel of your dashboards and widgets. For details, see the chapters Change the dashboard style (page 21) and Change the widget style (page 83).

8.5.9.2 Customize the application style

You can customize the look and feel of the entire application. Your changes of the application styling affect the MashZone NextGen welcome page (page 227), the dashboard editor in view and edit mode, and the feed editor. The styling of the Admin Console cannot be changed.

The styling information of the entire application are specified in the **application.less** file.

Prerequisites

You have MashZone NextGen administrator privileges.

Procedure

1. In database mode, perform the following steps.
 - a. Download the **application.less** file. (page 228)
 - b. Edit the styles in an appropriate text editor.
 - c. Save your settings.
 - d. Upload the modified file. (page 229)
2. In file-based mode, perform the following steps.
 - a. Edit the **application.less** style file located in the following folder.
<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\application
 - b. Compile the modified **application.less** file. (page 230)

Your changes are applied.

8.5.9.3 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 can change the style of the welcome page header.

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.

Prerequisite

You have administration privileges.

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\
Alternatively, you can specify a path to an image file using the **@brand-logo** parameter in the **application.less** file. The path must be absolute and begin with "/".
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 following folder.
<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\WEB-INF\config\
3. To change the appearance of the application header, you can edit the style templates supplied with MashZone NextGen. For details, see the chapter **Edit style templates**.
4. You can upload the customized **application.less** file in the Admin Console if you use the database mode. For details, see the chapter **Upload styles** (page 229).
5. If you use the file-based mode, perform the following steps.
 - a. Copy the customized **application.less** file in the following folder.
<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\application
 - b. Compile the modified application styles (page 230).

Your changes are applied.

MIGRATION OF THE APPLICATION.LESS FILE

If you have migrated the **application.less** file from MashZone NextGen version 9.12 or older, and you have customized 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;
```

8.5.9.4 Download styles

You can download available styles to use as style templates.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Styles & Themes** to expand this section of the **Administration** menu.

3. Click **Styles and Themes** to open the corresponding page.
4. To download the style file of the application, open the **Application** page and click **Download**.

The **application.less** style file is downloaded.

5. To download a dashboard or widget style file, open the **Dashboard** page and click the  **Download** icon of the desired style file.

The selected style file is downloaded and stored in the download folder specified in your web browser.

8.5.9.5 Upload styles

You can upload custom styles for your dashboards, widgets, and the entire application. The styles are stored in the MashZone NextGen database.

If you want to change or add a dashboard style, just upload the corresponding LESS style file, for example, myDashboardStyle.less. The LESS file can also be stored in a ZIP file. The ZIP file can have any name, for example, myStyles.zip.

If you want to change or add a dashboard style and widget styles, or just widget styles, you must upload a ZIP file that contains the corresponding LESS style files. The style files must be stored in the required folders, and the ZIP file must contain a corresponding file structure as described in chapter **How are the style files structured?** (page 233).

Prerequisite

You have administration privileges.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Styles & Themes** to expand this section of the **Administration** menu.
3. Click **Styles and Themes** to open the corresponding page.
4. To upload styles for the entire application, open the **Application** page.
 - a. Click **Upload**.
 - b. Select the **application.less** file to be uploaded.
 - c. Click **Open**.
5. To upload styles for a dashboard, open the **Dashboard** page.
 - a. Click **Upload**.
 - b. Select the **LESS** style file or the **ZIP** file that contains the style file to be uploaded.
 - c. Click **Open**.
6. To upload styles for a dashboard

The selected styles are uploaded to the MashZone NextGen database.

In the dashboard editor, you can select the uploaded styles to change the look and feel of your dashboards and widgets. For details, see the chapters Change the dashboard style (page 21) and Change the widget style (page 83).

8.5.9.6 Compile styles

If you use the file-based mode (page 231) to customize the styles of your dashboards and widgets, or the entire application, you must manually start the compiling of the modified style files to apply your changes in MashZone NextGen.

If you have modified widget styles, you must compile the corresponding dashboard style to apply your settings.

In database mode (page 231), the styles are automatically compiled when you upload the modified style file (page 229).

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
2. Click **Styles & Themes** to expand this section of the **Administration** menu.
3. Click **Styles and Themes** to open the corresponding page.
4. To compile application styles, open the **Application** page and click **Compile**.
5. To compile dashboard and widget styles, open the **Dashboard** page and click **Compile** for the desired styles and themes.

In file-based mode, the page lists all dashboard styles stored in the file system.

```
<MashZone NextGen installation\apache-  
tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\dashboard
```

You can only compile one style at a time. While compiling a style, all **Compile** buttons are disabled.

The selected styles are compiled.

CUSTOM WIDGETS

If you modify styles of a custom widget (page 233) while the server is started, you must manually start the compiling of the corresponding styles. Compile the dashboard styles that you want the custom styles to contain, such as the default dashboard style. In database mode, click the  Compile icon of the desired dashboard style.

8.5.9.7 Delete styles

You can delete dashboard and widget styles from the database.

Warning

Deleted styles and themes cannot be restored.

The application and the default styles cannot be deleted.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).

2. Click **Styles & Themes** to expand this section of the **Administration** menu.
3. Click **Styles and Themes** to open the corresponding page.
4. Open the **Dashboard** page and click the  **Delete** icon of the desired styles.

The selected styles are deleted from the database.

8.5.9.8 Switch the style editing mode

MashZone NextGen provides the database mode and the file-based mode to customize the application and dashboard styles. You can switch between the two editing modes.

DATABASE MODE

The database mode is the default mode to customize the look and feel of MashZone NextGen. All styles are stored in the MashZone NextGen database. A style can be downloaded from the database and edited to change it or create a new style. The changed style can then be uploaded back to the database to replace the old style or to add a new one.

FILE-BASED MODE

The file-based mode is used in a test or development environment to see immediate results of your customizing in MashZone NextGen. You can edit an existing style file, or add a new style file to your local installation. Then you must only compile your changed styles and the result of your changes will take effect immediately.

Warning

Switching the mode for a running system may cause unwanted effects. Instead of switching the mode for a running system, we recommend installing a test system to create and test style changes. The production system should always run in database mode.

Procedure

1. Open the **presto.config** file in an appropriate text editor. The file is located in the following directory.
`<MashZone NextGen installation>\apache-tomcat\webapps\mashzone\WEB-INF\classes\`
2. Specify the **enable.stylingconfiguration.fileBasedMode** parameter.
false: enables the database mode. This setting is default.
true: enables the file-based mode.
3. Save your changes.

Your settings are applied.

8.5.9.9 Adjust the font size to different devices

If you display your dashboards on devices with different screen sizes, for example, on mobiles, or tablets, you can adapt the font size of the displayed widgets.

The font size of most widget elements, such as the widget title, is automatically adjusted to the device size. But you can adjust the font size of all widget elements, such as the KPI values of a speedometer (see the example below).

By default, MashZone NextGen uses a predefined font size with the default size factor **1**. To adjust the font size to different devices, MashZone NextGen multiplies the default font size by the following predefined device factors:

Device	Device factor (Factors for media query)	Font size, for example
Desktop (default)	@default-factor: 1; min-width: 950px max-width: 1128px	12px
Mobile	@small-device-factor: 0.8; max-width: 529px	10px
Tablet	@medium-device-factor: 0.9; min-width: 530px) max-width: 949px	11px
Large devices	@large-device-factor: 1.2; min-width: 1129px	15px

The font sizes are specified in the **default.less** template file. The file is located in the following directory.

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

Example

You want to specify the font sizes for the KPI values of a speedometer. The font sizes for mobiles and tables are based on the predefined font size **font-big**. **Font-big** is set as default size for the widget element here.

```
@speedometer-kpi-value-font-size: @font-big;
@speedometer-kpi-value-font-size-mobile: ceil(@speedometer-kpi-value-font-size
* @small-device-factor);
@speedometer-kpi-value-font-size-tab: ceil(@speedometer-kpi-value-font-size
* @medium-device-factor);
@speedometer-kpi-value-font-array: @speedometer-kpi-value-font-size-
mobile @speedometer-kpi-value-font-size-tab;
```

Procedure

1. Open the **default.less** style file in your text editor.
2. Add your font size specifications for a widget element, as shown in the example above.
 - a. Set the default font size for your widget element.
For example:
@speedometer-kpi-value-font-size: @font-big;
 - b. Set the font size for different devices.

For example:

```
@speedometer-kpi-value-font-size-mobile: ceil(@speedometer-kpi-value-  
font-size * @small-device-factor);
```

```
@speedometer-kpi-value-font-size-tab:                ceil(@speedometer-kpi-  
value-font-size * @medium-device-factor);
```

- c. List the font sizes in a font array.

For example:

```
@speedometer-kpi-value-font-array: @speedometer-kpi-value-font-size-  
mobile                                @speedometer-kpi-  
value-font-size-tab;
```

3. Save your settings.
4. In database mode, upload the **default.less** (page 229).
5. In **file-based** mode replace the **default.less** file in the following directory and compile the modified style (page 230).

```
<MashZone NextGen installation>\apache-  
tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\dashboard
```

Your settings are applied.

8.5.9.10 Custom widgets

You can modify the styles of custom widgets (page 226), such as the Function Flow diagram.

Custom widgets are installed in the following folder:

```
<MashZone NextGen installation\apache-  
tomcat\webapps\mashzone\hub\dashboard\widgets\customWidgets
```

The style files of the custom widgets are automatically installed in corresponding subfolders of the **default** dashboard style.

Example

```
<MashZone NextGen installation\apache-  
tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-  
feel\dashboards\default\functionFlowWidget
```

The MashZone NextGen server automatically compiles the styles of newly installed custom widgets when the server starts.

Note that if you modify styles of a custom widget (page 233) while the server is started, you must manually start the compiling of the corresponding styles. In database mode (page 231), click the  **Compile** icon of the **default** dashboard style. On how to start the compiling of a dashboard style, see the chapter **Compile styles** (page 230).

8.5.9.11 How are the style files structured?

The different styles are specified in various **LESS** style files, such as application.less or default.less.

APPLICATION STYLING

The general styling settings of the entire application are specified in the **application.less** file. Changes in the **application.less** file affect the MashZone NextGen welcome page, the dashboard editor in view and edit mode, and the feed editor. The styling of the Admin Console cannot be changed.

The **application.less** file is stored in the following folder.

```
<MashZone NextGen installation>\apache-  
tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\application
```

DASHBOARD AND WIDGET STYLING

The dashboard and widget styles are specified in various style files with specific style names, such as White.less or Default-big.less. The default dashboard style file is named **default.less**.

The dashboard style files are store in the **dashboard** style folder:

```
<MashZone NextGen installation>\apache-  
tomcat\webapps\mashzone\hub\dashboard\assets\custom-look-and-feel\dashboard
```

Example

The LESS files of the **default** and **dark** dashboard styles:

```
...\assets\custom-look-and-feel\dashboard\  
▪ default.less
```

```
▪ dark.less
```

By default, a dashboard style provides a set of widget styles, that are stored in the corresponding widget style subfolder. The following folder names can be used to create the relevant widget subfolders: **actionbutton**, **barchart**, **bubblechart**, **colorpalette**, **columnchart**, **container**, **datefilter**, **grid**, **horizontalgauge**, **label**, **layoutgroup**, **layoutrow**, **linechart**, **piechart**, **speedometer**, **tab**, and **verticalgauge**.

Note that the list may vary depending on the MashZone NextGen version and the installed custom widgets.

Example

```
...\dashboard\default
```

```
  \actionbutton
```

```
  \barchart
```

```
  \bubblechart
```

```
  \colorpalette
```

```
  ...
```

For each widget style, a corresponding **LESS** style file is created and stored in the corresponding widget subfolder. The various styles must be assigned to the individual widget.

Example

The **Number_left_aligned** is assigned to the **Bar chart** and the **Action button** widget. The **Red** style is only assigned to the **Bar chart** widget.

...\dashboard\default\

- actionbutton\
Number_left_aligned.less
- barchart\
Number_left_aligned.less
Red.less

A widget style assigned to a dashboard applies to all widgets of the same type in this dashboard, for example, all bar charts in the dashboard.

FILE UPLOAD

You can upload your styles in individual LESS files or contained in a ZIP file. (page 229)

The structure of the ZIP file must correspond to the file and folder structure described above. The ZIP file can have any name.

Example

myStyling.zip

```
dashboardStyle.less
dashboardStyle
  actionbutton
    Default-big.less
    Primary-big.less
    Primary.less
  barchart
    White.less
  bubblechart
    White.less
  ...
  grid
```

8.5.10 Empty MashZone NextGen caches

You can manually delete the data of the internal caches that MashZone NextGen uses for caching the results of processing data sources and data feeds. By default, the caches are automatically emptied when the configured data refresh rate has expired.

MashZone NextGen provides a specific service to empty the cache on your local MashZone NextGen installation and the caches on all nodes in a multi-node environment remotely. You cannot empty the caches of individual nodes or a set of nodes in a multi-node environment.

Prerequisites

You have administration privileges.

Procedure

1. Start MashZone NextGen.
2. Log in to MashZone NextGen with administration privileges.
3. Enter the following URL in the address bar of your Web browser and press **Enter**.

`http://<hostname>:<port>/mashzone/mzservices/admin/clearresultcaches`

All caches have been emptied. A corresponding note informs you when the caches have been successfully emptied.

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 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 222), MashZone NextGen Event Service (page 222) and Command Central plug-in (page 141) for details concerning cluster configuration.

8.6.2 Use Events as Information Sources

To use events as information sources, you must create EDA or Apama event sources. Unlike other sources, only MashZone NextGen administrators can create EDA 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 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 (page 237)

- Create Apama Event Sources
- Start or Stop an Event Source (page 258)

You can also Manage Apama Event Sources (page 247) and Manage EDA Event Sources (page 237).

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 (page 237)
 - Edit EDA Event Sources (page 241)
 - Duplicate EDA Event Sources (page 245)
 - Delete EDA Event Sources (page 246)
 - Share EDA Event Sources (page 246)

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 106).
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	Require d	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		Available only if Strategy is set to Buffer . 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. <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	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.
Consider dimension		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 .
Dimension attribute	conditional	Available only when Strategy is set to BUFFER and required when the Consider dimension option is set. Select the event attribute whose unique values determine separate event series (buckets) for this event source.
Max. number of dimension values	conditional	Available only when Strategy is set to BUFFER and required when the Consider dimension option is set. 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. Default value: 1 (Max: 100.000). The product of Max. number of dimension values and Capacity per dimension value must not be greater than 100.000.
Dimension Squeeze-out		Available only when Strategy is set to BUFFER and required when the Consider dimension option is set. 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. This option is clear by default which discards new events with new unique dimension values once the maximum number of buckets has been reached. 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). Default value: false
Capacity per dimension value	conditional	Available only when Strategy is set to BUFFER and required when the Consider dimension option is set. Enter the maximum number of events that can be stored in a specific event series (bucket) for each unique dimension value. Default value: 10 The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.
Event ID attribute	yes	Available only when Strategy is set to DELTA .

		Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.
Command attribute	yes	Available only when Strategy if set to DELTA . 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.
Capacity	yes	Enter the maximum number of events to store in this event source. (Max: 100.000) Default value: 10
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 Max: stores events in a local BigMemory Max cache.
Throttling		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: <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		Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1): <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 106).
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		Available only if Strategy is set to Buffer . 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. <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	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.</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>

		The product from Max number of dimension values and Capacity per dimension value must not be more than 100 000.
Event ID attribute	yes	Available only when Strategy is set to DELTA . Select the attribute that identifies an event. The event ID and command determines which events are stored, updated or removed in this event source.
Command attribute	yes	Available only when Strategy if set to DELTA . 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.
Capacity	yes	Enter the maximum number of events to store in this event source. (Max: 100.000) Default value: 10
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 Max: stores events in a local BigMemory Max cache.
Throttling		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: <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		Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1): <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 (page 237)

8.6.3.3 Duplicate EDA Event Sources

You can duplicate existing EDA Event Source.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 106).
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 106).
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. 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.
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.

8.6.4 Manage Apama Event Sources

You can create, edit, delete, import, and share **Apama** Event Sources:

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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
 - Edit Apama Event Sources
 - Duplicate Apama Event Sources (page 252)
 - Delete Apama Event Sources (page 252)
 - Share Apama Event Sources (page 258)

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 106).
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 the 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 the 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 want 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 (page 259) 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 you can select a scenario to which this event source should subscribe.
Strategy	yes	<p>The strategy that this event source uses for saving and removing events published by the Event Bus. Valid strategies are:</p> <ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until the 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 the 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 cannot contain all possible fields for the event. Simply put, each event can 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 the event source memory is full. When 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. MashZone NextGen automatically determines the unique identifier of each event, defined by one or more key fields. MashZone NextGen maintains a single row for each event key, updating its values by processing the Apama scenario InstanceUpdated Events. 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. 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.
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. 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. 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. Default value: 10 (Max: 100.000). The result of multiplying Max. number of dimension values by Capacity per dimension value must not exceed 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set. Determines how additional events are handled when they have new unique values for the dimension that defines</p>

		<p>buckets in this event source, but the maximum number of unique values (buckets) has already been reached.</p> <p>This option is not set by default. the option discards new events with new unique dimension values when 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 result of multiplying Max number of dimension values by Capacity per dimension value must not exceed 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 determine 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	Enter the maximum number of events to store in this event source. (Max: 100.000) Default value: 10
Memory model		Determines where events are stored: <ul style="list-style-type: none"> ▪ Internal: default. Stores events in local memory for this event source. ▪ BigMemory Max: stores events in a local BigMemory Max cache.
Throttling		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 do the following: <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 is received and change the number, if needed. See Example below.
Exception		Set this option to support a hybrid throttling strategy, typically involving both time and event limitations. Then set the exception criteria (Default=1): <ul style="list-style-type: none"> ▪ A number ▪ Milliseconds or Events as the measurement for the exception criteria See Example below.

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 subscribed real-time views. Throttling can slow event updates to real-time views when the volume or frequency of 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 as soon as a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies might not be able to balance the 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:

With the example hybrid throttling strategy shown above:

- The event source waits 50 milliseconds after pushing events to subscribed views.
- If fewer than 10 events are received in these 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 the subscribed views, and both the time interval and the count of events begin 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 subscribed 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 106).
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 106).
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 106).
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 the 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	Enter a unique name for this event source.
Start the 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 want 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 (page 259) 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 you can select a scenario to which this event source should subscribe.
Strategy	yes	The strategy that this event source uses for saving and removing events published by the Event Bus. Valid strategies are:

	<ul style="list-style-type: none"> ▪ BUFFER = FIFO (first in-first out). Events are stored until the 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 the 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 cannot contain all possible fields for the event. Simply put, each event can 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 the event source memory is full. When the event source memory is full, new events are discarded.</p> ▪ 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. MashZone NextGen 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. 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. 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. 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. Default value: 10 (Max: 100.000). The result of multiplying Max. number of dimension values by Capacity per dimension value must not exceed 100.000.</p>
Dimension Squeeze-out		<p>Available only when Strategy is set to BUFFER and required when the Consider dimension option is set. Determines how additional events are handled when 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. This option is not set by default. the option discards new events with new unique dimension values when the maximum number of buckets has been reached. 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). 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 result of multiplying Max number of dimension values by Capacity per dimension value must not exceed 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 determine 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: default. Stores events in local memory for this event source. ▪ BigMemory Max: stores events in a local BigMemory Max 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 do the following:</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 is 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 subscribed real-time views. Throttling can slow event updates to real-time views when the volume or frequency of 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 as soon as a minimum count of events are received, such as 10 events. For example:

Throttling * Exception...

-- or --

Throttling * Exception...

Simple strategies might not be able to balance the 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 waits 50 milliseconds after pushing events to subscribed views.
- If fewer than 10 events are received in these 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 the subscribed views, and both the time interval and the count of events begin 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 subscribed 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 106).
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. 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. 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.

8.6.5 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

1. Open the MashZone NextGen Admin Console (page 106).
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 event source tab **EDA** 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.

8.6.6 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 106).
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.

8.6.7 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 106).
2. Click **Event Service** to expand this section of the **Administration** menu.
3. Click **Apama instances**.
4. Select further steps:

Create Apama Instances (page 260)

Edit Apama Instances (page 260)

Delete Apama Instances (page 261)

8.6.7.1 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 106).
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	Require	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)

8.6.7.2 Edit Apama Instances

You can edit an already existing Apama instances.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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	Require	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)
------	-----	--

8.6.7.3 Delete Apama Instances

You can delete Apama instances.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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.

8.6.8 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 106).
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 (page 261)
 - Edit Apama Event Targets (page 262)
 - Delete Apama Event Targets (page 263)
 - Share Apama Event Target (page 263)

8.6.8.1 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 106).
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 (page 259) 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.

8.6.8.2 Edit Apama Event Targets

You can edit an already existing Apama event target.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 (page 259) 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.

8.6.8.3 Delete Apama Event Targets

You can delete Apama event targets.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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.

8.6.8.4 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 106).
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. 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. 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.

9. 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 106).
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 264)
 - Edit PPM Connections (page 265)
 - Delete PPM Connections (page 266)
 - Share PPM connections (page 267)

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 online help topics.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 using 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 online help topics.

Procedure

1. Open the MashZone NextGen Admin Console (page 106).
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 using 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 106).

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 106).
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. 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.

6. 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.

7. 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 110) 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|reset_password)\.html.*" security="none" request-
matcher="regex"/>
<http pattern="/help/.*" 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="jwTokenProcessingFilter" 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>
```

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 79) 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 157) 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 124)
- Support International Character Sets and Locales (page 154)
- Use the Default MashZone NextGen User Repository (page 138)
- Change MashZone NextGen Repository Ports (page 156)
- Tuning the MashZone NextGen Repository Connection Pool (page 271)
- Synchronize the MashZone NextGen Repository and MashZone NextGen Server Time Zones (page 272)
- Sharing the MashZone NextGen Repository in Clustered Environments (page 217)
- Configure BigMemory Max Servers for MashZone NextGen Caching and In-Memory Stores (page 170)
- Maintenance Suggestions (page 271)

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 the official Tomcat Datasource Properties.

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 124) 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 37).

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 .
Show values	Displays the measure values outside the bars by default. If the space is too limited, the values are displayed inside the bars. If the space inside the bars or stack bars is also limited, the values are not displayed at all.
Legend position	Displays a legend in the widget and sets the legend's position. The default is None , that is, no legend is displayed.

Specific bar chart options	Description
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>
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 (page 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 (page 74) for details.</p>

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.

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 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 47)
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.

General options	Description
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.

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.

General options	Description
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 column chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
X-axis label orientation	<p>Aligns the labeling of the X-axis.</p> <ul style="list-style-type: none"> ▪ Auto: aligns the labels depending on the available space on the X-axis. This behavior is default. ▪ Rotated: rotates the labels of the X-axis 90 degrees. ▪ Alternate: the labels of the X-axis are alternately offset.
Show values	Displays the measure values outside the columns by default. If the space is too limited, the values are displayed inside the columns. If the space inside the bars or stack columns is also limited, the values are not displayed at all.
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.

Specific column chart options	Description
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 (page 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 (page 74) 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	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 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. 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 From and To values of the default selection are set to the same value.</p> <p>If you are using a PPM version older than 10.5, this option is not available for Process Mining context-based dashboards.</p>
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	Specifies the range values that are preselected by default.

Specific date filter options	Description
	<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 Distribution chart

You can use a distribution chart to identify outliers and suspected deviations for process or function instances.

In a distribution chart, you can display the frequency distribution of a distributable measure's values for a number of process or function instances.

The distribution chart is available only for context based dashboards.

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.

General options	Description
	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 distribution chart options	Description
Chart title	Optional chart title of the chart displayed in the widget.
X-axis label orientation	Aligns the labeling of the X-axis. <ul style="list-style-type: none"> ▪ Auto: aligns the labels depending on the available space on the X-axis. This behavior is default. ▪ Rotated: rotates the labels of the X-axis 90 degrees. ▪ Alternate: the labels of the X-axis are alternately offset.
Show values	Displays the measure values outside the columns by default. If the space is too limited, the values are displayed inside the columns. If the space inside the bars or stack columns is also limited, the values are not displayed at all.
Axis	Specifies the visualization of the X- and Y-axis. <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>

9.1.8 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.
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>

General options	Description
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.

9.1.9 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.
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>

General options	Description
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	<p>Enables you to edit the Grid widget interactively. Click the Configuration icon. In edit mode, the Grid widget has a blue frame.</p> <p>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 (page 9).</p>
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 (page 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 (page 74) for details.

9.1.10 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.

Specific circular gauge chart options	Description
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.
Scale size	Selects the scale size of the gauge chart. The default is Medium .

9.1.11 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 78).

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.

9.1.12 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 .

9.1.13 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 .

9.1.14 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.

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 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.
X-axis label orientation	Aligns the labeling of the X-axis. <ul style="list-style-type: none"> ▪ Auto: aligns the labels depending on the available space on the X-axis. This behavior is default. ▪ Rotated: rotates the labels of the X-axis 90 degrees. ▪ Alternate: the labels of the X-axis are alternately offset.
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. <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>

9.1.15 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.
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 (page 9).

Specific list options	Description
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 (page 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 (page 74) for details.</p> <p>The option is enabled by default.</p>

9.1.16 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 Upload tile servers.</p>
Default marker icon	<p>Selects the default marker icons.</p> <p>Default icon is a circle.</p> <p>Select an icon in the drop-down menu.</p>
Sticky active area	<p>Adjust automatically the zoom level of the map based on the active markers to display all markers in the widget.</p>
Show label	<p>Displays the label of all regions if applicable.</p>

9.1.17 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>

General options	Description
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 (page 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 (page 74) for details.

9.1.18 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 31). 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 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 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

Specific Rich text area options	Description
	<ul style="list-style-type: none"> ▪ 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>

9.1.19 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.

General options	Description
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 maximum (To) value of a data range. ▪ 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. ▪ Range (open): With this slider, you can select the minimum (From) value and the maximum (To) value of a data range. You can also specify an open range, that means, you can set a fixed minimum and an open maximum or an open minimum and a fixed maximum. <p>When you move the slider beyond the start and end point of the slider, the selection is removed.</p>
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>

Specific slider options	Description
	<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>For the Range (open) type, you specify the From and 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>
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>

Specific slider options	Description
Rotate axis	Displays the values at an angle of 45° on the axis if the option is enabled, otherwise displays the values horizontally. The option is disabled by default.
Numeric format	Selects the format of the values displayed in the widget. The slider widget supports only numeric values.

9.1.20 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.
Container	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. Hide border: Hides the outline of the component container. Click the icon again to display the outline. The border is visible by default.
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.

Specific Jump to PPM widget options	Description
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.).

9.1.21 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	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.

General options	Description
	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 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 .

9.1.22 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	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.

General options	Description
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	Selects a map template of outline maps based on available GeoJSON configurations. Default map template is world-countries-by-name. Select a map template in the drop-down menu. You can add your own map templates. For details, see Upload geographical maps.
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. 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 303)
A data source operator enables you to specify the connection to a data source and to configure the data retrieval.
- Data transformation operators (page 323)

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 360)

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. 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 200). <p>The file size is limited to 80 MB.</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>

Parameters	Description
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 Max

Extracts data from a BigMemory Max cache.

To use the BigMemory Max source operator, you have to configure a BigMemory Max connection in MashZone NextGen in advance, see BigMemory Max for Caching, Connections and In-Memory Stores (page 168).

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: <code>SELECT * FROM</code> . See example BigMemory Max 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 (page 322) for details.
Expand ()/ Collapse ()	Expands or collapses the query input box.


Parameters	Description
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.

For detailed information on configuring the time zone used by this operator, see the chapter Configure feed processing time zone (page 153).

Example: BigMemory Max data source

To use the RAQL source operator, you have to configure a BigMemory Max connection in MashZone NextGen in advance, see BigMemory Max for Caching, Connections and MashZone NextGenMashZone NextGen (page 168).

To use the RAQL BigMemory Max 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	<p>Text file, with values that are separated by the same separator.</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 200). ▪ File alias: 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> <p>The file size is limited to 80 MB.</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> <p>For details, see URL parameter syntax (page 321).</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
Authentication	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.
HTTP headers	Adds HTTP headers to the URL.
Separator	Separates the column values in the CSV file. Comma (,), semicolon (;), space, tab and pipe () are available. Default is comma.
Get column names from row	Specifies a specific row that contains the column names. Activate the Get column names from row option and enter the number of the relevant row.
Data from row	Specifies a specific row from which the data source values extraction starts.
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	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.

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 196). 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.

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 96) 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.

Parameters	Description
	For details, see Event Service Configuration and Administration (page 236).
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.

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 182).

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 (page 322).
Expand ()/ Collapse ()	Expands or collapses the query input box.

Parameters	Description
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 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 200). File alias: 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 <p>The file size is limited to 80 MB.</p>

Parameters	Description
Browse file alias (...)	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 198).
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 321).
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, <code>?catalog?journal?articles?*</code> <code>?data?rows?*</code> <code>?catalog?books?*</code></p> <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>

Parameters	Description
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 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.

For detailed information on configuring the time zone used by this operator, see the chapter [Configure feed processing time zone](#) (page 153).

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 196).

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.

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 200). ▪ File alias: 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> <p>The file size is limited to 3 MB.</p>
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 198).</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 321).</p>

Parameters	Description
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.
Sheet	Sheet in the source table to be extracted. The default value: First sheet Specification: Mandatory
List table / Cross table	Specifies the table type. The default value: List table Specification: Mandatory For cross tables, only a single vertical iteration on the left side of the table is currently supported.
Separator	Separates the column values in the CSV file. Comma (,), semicolon (;), space, tab and pipe () are available. Default is comma.
Column name from row	Determines the names of the individual columns from a specific row, the row number of which must be specified. This option is not available for cross tables.
Horizontal iteration from row	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.
Import values from row	Extracts all values from the source file starting with a specific row. The default value: 2 Specification: Mandatory
Import data range from/to	Area of the table from which data is to be extracted, specified using column and row coordinates, for example, A3 to H128 Specification: Optional You can only specify a single continuous data range that can contain empty rows or columns. If no upper limit (to) is specified for the data range, all cells above the lower limit (from) are extracted.
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

Parameters	Description
	reloaded from the data source and all changes in the list are undone.

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 196).

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.

9.2.1.9 PPM

Uses the ARIS Process Performance Manager (PPM) query interface to retrieve data from favorites defined in PPM.

For every PPM system that can be used in MashZone NextGen, a PPM connection (page 264) 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.

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. For details, see Process Performance Manager Integration (page 264).
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 URL and insert the favorite URL created in PPM.

Parameters	Description
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 Authentication with single sign-on solutions (page 110).</p> <p>HTTP basic auth: Requires the user name and the password of a PPM user.</p> <p>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	<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>
Filters: Detect	<p>Reloads the dimensions and measures information from PPM</p>
Add filter	<p>Adds one or more filter criteria to filter the values of the PPM query.</p> <p>PPM supports filtering by multiple selection. For details on using multiple selections, see Use multiple selection in lists and tables (page 9).</p> <p>For details on configuring the PPM operator for filtering, see Add filters (page 316).</p>

9.2.1.9.1 Add filters

You can add one or multiple filter criteria to filter the values of the PPM query.

The **PPM** data source operator supports user inputs of **List** type as input parameter. Thus, widgets using the **PPM** operator allows the use of multiple selection (page 9). That means that you can select one or multiple values in another widget to filter values provided by the **PPM** data source operator. To enable multiple selection when using the **PPM** operator, you must assign at least one user input of **List** type as input parameter. You can select **Date user input (List)** and **Text user input (List)** as input parameter. **Number user input (List)** is currently not supported by the **PPM** operator.

For details on creating input parameters, see the chapter Create input parameters (page 70).

Procedure

1. In the **PPM** operator, click **Add filter**.
2. Select a dimension or measure, such as **Date** or **Process cycle time**, and click **Add**. The selected criterion is added to the PPM operator as a filter criterion.
3. Select an operator for the condition of the filter criterion in the drop-down menu, for example, **is equal to**.

For criteria of **Text** type, the **is equal to** operator is available by default. If you enable filtering by expressions, the **starts with**, **ends with**, and **contains** operators are additionally provided. For enabling filtering by expressions, see step **5** below.

For criteria of **Date** type, the operators **On**, **Before or on**, and **On or after** operators are available. The **Before or on** and **On or after** operators cannot be used with **Date user input (List)**.

4. Enter a constant filter value in the input field, or click  **Insert input parameter** and select an input parameter. The button can only be clicked if there is at least one input parameter with the same data type as the filter criterion.

To enable the multiple selection for the **PPM** operator, select a user input of **List** type as input parameter.

5. You can filter criteria of **Text** type by expressions instead of entire words. For example, if you use a place holder in the filter expression, for example, * or ?.
 - a. Click the  **Set filter properties** icon and specify the filter parameters.
 - b. Enable the **Filter using an expression** uses.
 - c. Select whether the **Key** or the **Description** of the criterion is to be used as filter value.

If you enable filtering by expressions, further operators are provided to set the filter condition. See step **3** above.

6. Click **Add condition (+)** to add further filter conditions.

If you add further filter condition, the **AND-link conditions** option is displayed. You can enable the option to link the filter condition with a logical AND. By default, the conditions are linked with a logical OR.

7. Click **Delete condition (-)** to remove a filter condition.

Your settings are applied.

9.2.1.10 Terracotta

Extracts data from datasets managed by a Terracotta 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 97). You must select one of those two categories of datasets. Depending on your selection, the source operator shows different parameters.

To use the Terracotta source operator, you must configure a Terracotta connection in MashZone NextGen in advance. Terracotta connections can be defined in the MashZone NextGen **Admin console**. See Manage Terracotta connections.

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 alias	<p>List of available Terracotta datasets for the configured connection aliases.</p> <p>Only Dataset aliases for which you have the Usage privilege are available. Select a Dataset alias. See Manage Terracotta connections 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 which you have the Usage privilege are available. Select the alias that 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 190) for details.</p> <p>The option is available if the Materialized Feeds source type is selected.</p>
Dataset name	<p>List of available Terracotta datasets for the configured connection aliases. 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 in the Dataset Alias drop-down box.</p>
Advanced options	<p>Specifies further parameters.</p>

Parameters	Description
Batch size	Stream executions attempt (if possible) to optimize execution by transferring multiple elements simultaneously over the network. The Batch size option specifies the batch sizes to be used when transferring elements over the network. The Terracotta server can limit the actual value. For more information, please refer the corresponding Terracotta documentation.
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 (page 322).
Expand ()/ Collapse ()	Expands or collapses the query input box.
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 you select 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. Use Configure columns to display or modify that schema.</p> <p>Disable 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 must match the cell definition in the store. If the definition does not fit the definition, the corresponding values are always null. For example, if the cell definition is age with type INT and you specify age with type DOUBLE, the column definition does not match the cell definition and the values are null.</p>

For detailed information on configuring the time zone used by this operator, see the chapter [Configure feed processing time zone \(page 153\)](#).

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: 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 200). ▪ File alias: 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> The file size is limited to 80 MB.
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	Description
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> <p>/data/row</p> <p>/catalog/book</p>
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.

For detailed information on configuring the time zone used by this operator, see the chapter [Configure feed processing time zone](#) (page 153).

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 Max**, **JDBC**, **Terracotta**, 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.

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.

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.

PARAMETERS

The following parameters are available.

Action	Result
Dimension column	Name of dimension column. Source: Source table

Action	Result
	Data type: Date, Number, or Text The default value: {None} Specification: Optional
Aggregation column	Name of the column to be aggregated. Source: Aggregation column is transferred from the source table. Data type: Number The default value: {None} Specification: Optional
Aggregation type	Aggregation type for the column to be aggregated: Average value, Sum, Minimum, Maximum, Count, First row and Last row. 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

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.

Data type	Comparison operators
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.

EXAMPLE

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. 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)

Action	Result
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, cos, tan, asin, acos, and atan in degrees.</p>

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.

The following parameters are available.

Parameter	Description
Column	<p>Name of the column for which the average value is calculated. Column is transferred from the source table.</p> <p>Data type: Number</p> <p>Specification: Mandatory</p>
Weight	<p>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.</p> <p>Data type: Number</p> <p>Specification: Optional</p>
Target column	<p>Name of the column to which the result is written. The column name can be transferred from the source table or freely entered.</p>

Parameter	Description
	Data type: Number Default value: Result_1 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	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; 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.

Action	Result
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 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.

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>

Parameter	Description
	<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>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>

Parameter	Description
Specify type	Automatically specifies the data type of the source values. 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.
Encoding (Text source format)	Specifies the encoding of special characters, for the conversion of the Text type to the Text type, for example, "/", "&", "?". Apply UTF-8 encoding: encodes the entire text, using UTF-8 codes Decode UTF-8: decodes the entire text, using UTF-8 codes Apply URL encoding: keeps the URL specific characters 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.

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".

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; separated by separators. In addition, the

Action	Result
	<p>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	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.

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".

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

Converts a data feed column into a value or value list. If the **to list** option is enabled, the resulting value is a list containing values from all rows of the column. Otherwise, the first value found is returned.

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, depending on the input column type; Specification: Mandatory
to list	Creates a value list. If enabled, the resulting value is a list of all column values. Otherwise the first value found is returned. Specification: Optional

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.

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.

Options	
Include key values of left data feed	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

	<p>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.

The following parameters are available.

Parameter	Description
Left/right data feeds	<p>Two data feeds to be combined.</p> <p>Specification: Mandatory</p>

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.

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".

The following parameters are available.

Action	Result
Column	Name of the column whose values are replaced. Source: Source table Data type: Number, Text, Date; Specification: Mandatory
New value	Value that replaces the value in the source column. Source: Column with values, single value from a feed (single-value operator), user input, or a constant. Default value: is equal to Specification: Mandatory
Replace	Values are replaced if one or all conditions is/are met.
Source column	Name of the column whose values are compared. Source: Source table Data type: Number, Text, Date; Specification: Mandatory
Comparison operator	Operator that compares the values from the source column with the comparison values.

Action	Result
	<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>

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

Data type	Comparison operators
Date	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. The following parameters are available.

Parameter	Description
Column	Name of the column whose values are converted. Source: Source table Data type: Text Specification: Mandatory
Conversion	Transformation rule for conversion of column values: UPPER: Converts all characters into upper case, according to the rules of the specified language. LOWER: Converts all characters into lower case, according to the rules of the specified language. ONLY_LETTERS: Removes all figures (0-9) from the column values; ONLY_NUMBERS: Removes all letters from the column values. REMOVE_SPACES: Removes all spaces from the column values. REMOVE_LEADING_WHITESPACE: Removes leading whitespace from the column values. REMOVE_TRAILING_WHITESPACE: Removes trailing whitespace from the column values. Specification: Mandatory
Target column	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). Data type: Text

Parameter	Description
	Default value: Result_1 Specification: Optional If the target column is identical to the source column, the values in the source column are overwritten.

9.2.2.12 Copy data feeds

Creates up to four independent copies of a data feed.

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 value or list value, without changing the input value.

The following parameters are available.

Parameter	Description
Input value	Input value or value list to be copied. Specification: Mandatory

9.2.2.14 Delete column

Deletes the specified columns from the data feed.

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.

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.

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. Source: Source column, single-value operator, input value, or constant. Size ≥ 0 ; Data type: Number Specification: Mandatory

Parameter	Description
Length	<p>Number of characters in the string to be extracted.</p> <p>Source: Source column, single-value operator, user input, or constant.</p> <p>Number of characters ≥ 0</p> <p>Data type: Number</p> <p>Specification: Mandatory</p>
Target column	<p>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).</p> <p>Data type: Number</p> <p>Default value: Result_1</p> <p>Specification: Mandatory</p>

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.

The following parameters are available.

Parameter	Description
Source column	<p>Name of the source column for which the earliest or latest date values are determined.</p> <p>Source: Source table</p> <p>Data type: Date</p> <p>Specification: Mandatory</p>
Earliest/latest date	<p>Determines the earliest or latest date values in the source column.</p> <p>Default value: Earliest date</p>
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 360).

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>
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>

Action	Result
All conditions should match.	Combines the specified filter conditions. If the option is not enabled, the result will be a concatenation of all single filter results.

The following parameters are available.

Parameter	Description
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 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.

The following parameters operators 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
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%.

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 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.

Parameter	Description
	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.

The following parameters are available.

Action	Result
Column name	Name of the new column. 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.

The following parameters are available.

Parameter	Description
Text	Any strings 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.

For 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.

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 Specification: Mandatory
Value	Value by which the date is moved by the selected unit.

Action	Result
	<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>

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 Named copy

Creates an independent copy of a value or list value and assigns a name to it. The input value is not changed.

The following parameters are available.

Parameter	Description
Input value	Input value or value list to be copied. Specification: Mandatory
Name	The name assigned to the value. Default: Copy

9.2.2.25 RAQL Inline

Transforms up to four data feeds into one output by applying a RAQL statement.

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 (page 322) 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 370).

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.

For detailed information on configuring the time zone used by this operator, see the chapter Configure feed processing time zone (page 153).

EXAMPLE

```
SELECT feed1.productName, feed2.supplierName FROM feed1, feed2 WHERE price <= :maxPrice AND feed1.supplierId = feed2.id
```

9.2.2.26 Rename column

Changes the names of the specified columns from the data feed. The data type of the column is retained.

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

9.2.2.27 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.

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 a empty text.
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

Parameter	Description
	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".

9.2.2.28 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.

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 transferred from the source table or freely entered. Data type: Number Default value: Result_1 Specification: Optional

9.2.2.29 Round up/down date

Converts date values from a date column to a rougher time unit and writes the results to a target column.

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>
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>

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	2010-08-06T 17 :00
2010-08-06T17: 18 :00	15 minutes	2010-08-06T 17:15
2010-08-06T 02:18:04	4 hours	2010-08-06T 04:00:00

Round up

Source value	Accuracy	Result
2010-02-28T23: 07:00:00 AM	15 minutes	2010-02-28T 23:00
2010-02-28T23: 07:30	15 minutes	2010-02-28T 11:15:00 PM
2010- 02-28 T23:30:00	1 hour	2010- 03-01 T00

Round down

Source value	Accuracy	Result
1970:01:01T09:00: 01	6 hours	1970:01:01T 12:00:00
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- 02-28 T 20:00:01	6 hours	2010- 03-01 T00

Always round down

Source value	Accuracy	Result
2010-02-28T 11:15:00 PM	15 minutes	2010-02-28T 11:15:00 PM

Source value	Accuracy	Result
2010-02-28T 11:14:59 PM	15 minutes	2010-02-28T 23:00
2010- 03-01T05:59:59	6 hours	2010- 03-01T00

9.2.2.30 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.

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 Specification: Mandatory</p>
Property	<p>Properties of the logged-in user; displayed if User data is selected as the information type.</p> <p>The following values can be selected: Login, First name, Last name, E-mail and Language</p> <p>Default value: User name 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 Specification: Mandatory</p> <p>For detailed information on configuring the time zone used by this operator, see the chapter Configure feed processing time zone (page 153).</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 Specification: Mandatory</p>

Parameter	Description
Upper/lower limit	Upper or lower limit of the value range for the random number; displayed if Random number is selected as the information type. Default value: 0 and 10 Specification: Mandatory

9.2.2.31 Value to column

Converts an individual value into a column so that it can be connected to an operator.

Creates a feed table from an input value. If the input value is a list, the result contains one list element per row. Otherwise the result has a single row containing the value.

The following parameters are available.

Parameter	Description
Target column	Name of the column to which the conversion result is written. The name of the column is initially Result . You can change the name manually if required. Data type: Text, Date or Number, depending on the input value. Default value: Result 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.

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: 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

The date user input has the format yyyy-MM-dd'T'HH:mm:ss (up to the required accuracy) or yyyy-'Q'Q.

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.
Default value	Value used for a test calculation in the feed editor.

Parameter	Description
	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).

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

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.
Default 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.

The following parameters are available.

Parameter	Description
Name	Name of user input

Parameter	Description
	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 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

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.
Default value	Value used for a test calculation in the feed editor.

Parameter	Description
	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
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 (page 354).

In feed editor and dashboard editor you can select the BigMemory Max operator and run a RAQL query against data available in a BigMemory Max server. See also BigMemory Max (page 304).


In feed editor and dashboard editor you can select the Terracotta operator and run a RAQL query against data available in a Terracotta server. See also Terracotta (page 317).

RAQL queries are processed within a dedicated analytical engine, which is an integral part of the MashZone NextGen server.

9.3.1.3 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 368) 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.

9.3.1.4 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
float	java.lang.Float	number
integer	java.lang.Integer	number
long	java.lang.Long	number
short	java.lang.Short	number

RAQL Datatype	Java Equivalent	Query Result Mapping
string	java.lang.String	string

9.3.2 Language reference

9.3.2.1 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			

9.3.2.2 Syntax Reference

This topic contains syntax definitions and corresponding syntax diagrams for the Real-Time Analytics Query Language (RAQL) in MashZone NextGen.

9.3.2.2.1 ExplainOrSelectStatement



Definition:

ExplainOrSelectStatement ::= ('EXPLAIN' ('(' CharacterConstant ')')?)? SelectStatement

Used In:

This is the starting definition of RAQL queries, which has no parent references.

9.3.2.2.2 CharacterConstant




Definition:

CharacterConstant ::= S_CHAR_LITERAL | S_PARAMETER

Used In:

- ExplainOrSelectStatement (page 370)

9.3.2.2.3 SelectStatement



Definition:

SelectStatement ::= SelectWithOrder ';'?

Used In:

- ExplainOrSelectStatement (page 370)

9.3.2.2.4 SelectWithOrder




Definition:

SelectWithOrder ::= QueryExpressionBody OrderByClause? LimitClause?

Used In:

- SelectStatement (page 370)
- SourceReferenceOrSubquery (page 392)

9.3.2.2.5 QueryExpressionBody




Definition:

QueryExpressionBody ::= QueryTerm (('UNION' | 'EXCEPT' | 'MINUS') ('ALL' | 'DISTINCT')? QueryTerm)*

Used In:

- QueryPrimary (page 372)
- SelectWithOrder (page 371)
- SqlRelationalOperatorExpression (page 378)

9.3.2.2.6 QueryTerm




Definition:

QueryTerm ::= QueryPrimary ('INTERSECT' ('ALL' | 'DISTINCT')? QueryPrimary)*

Used In:

- QueryExpressionBody (page 371)

9.3.2.2.7 QueryPrimary




Definition:

QueryPrimary ::= SelectWithoutOrder | '(' QueryExpressionBody .'

Used In:

- QueryTerm (page 371)

9.3.2.2.8 SelectWithoutOrder




Definition:

SelectWithoutOrder ::= 'SELECT' ('ALL' | 'DISTINCT')? SelectList FromClause WhereClause? GroupByClause? HavingClause?

Used In:

- QueryPrimary (page 372)

9.3.2.2.9 SelectList




Definition:

SelectList ::= SelectItem (',' SelectItem)*

Used In:

- SelectWithoutOrder (page 372)

9.3.2.2.10 SelectItem




Definition:

SelectItem ::= '*' | SelectStar | SqlSimpleExpression Alias?

Used In:

- SelectList (page 372)

9.3.2.2.11 SelectStar




Definition:

SelectStar ::= ObjectName '.' (ObjectName '.')? '*'

Used In:

- SelectItem (page 373)

9.3.2.2.12 ObjectName




Definition:

ObjectName ::= S_IDENTIFIER | S_QUOTED_IDENTIFIER

Used In:

- Alias (page 392)
- GenericFunctionCall (page 387)
- SelectStar (page 373)
- SourceField (page 387)
- SourceReference (page 393)
- SqlPrimaryExpression (page 376)

9.3.2.2.13 SqlSimpleExpression



Definition:

**SqlSimpleExpression ::= SqlMultiplicativeExpression (('+' | '-' | '||')
SqlMultiplicativeExpression)***

Used In:

- BooleanCaseCall (page 385)
- SelectItem (page 373)
- SimpleCaseCall (page 385)
- SqlBetweenClause (page 380)
- SqlLikeClause (page 380)
- SqlRelationalExpression (page 378)
- SqlRelationalOperatorExpression (page 378)

9.3.2.2.14 SqlMultiplicativeExpression




Definition:

**SqlMultiplicativeExpression ::= SqlExponentExpression (('*' | '/' | '%')
SqlExponentExpression)***

Used In:

- SqlSimpleExpression (page 374)

9.3.2.2.15 SqlExponentExpression




Definition:

SqlExponentExpression ::= SqlUnaryExpression ('' SqlUnaryExpression)***

Used In:

- `SqlMultiplicativeExpression` (page 374)

9.3.2.2.16 SqlUnaryExpression




Definition:

SqlUnaryExpression ::= ('+' | '-')? SqlPrimaryExpression

Used In:

- `SqlExponentExpression` (page 375)

9.3.2.2.17 SqlPrimaryExpression



Definition:

```
SqlPrimaryExpression ::= '(' SqlExpression ')' | Constant | FunctionCall | (
AggregateFunctionCall | GenericFunctionCall ) WindowSpecification? |
GroupingOperation | WindowFunctionCall | ObjectName ( '.' ( GenericFunctionCall
WindowSpecification? | ObjectName ( '.' ObjectName )? ArrayElementReference? ) |
ArrayElementReference )? | ArrayValueConstructor
```

Used In:

- SqlUnaryExpression (page 375)

9.3.2.2.18 SqlExpression



Definition:

SqlExpression ::= SqlAndExpression ('OR' SqlAndExpression)*

Used In:

- AggregateFunctionCall (page 386)
- BooleanCaseCall (page 385)
- FunctionCall (page 382)

- GroupingExpressionReference (page 396)
- HavingClause (page 398)
- JoinedTable (page 393)
- OrderByClause (page 388)
- SqlExpressionList (page 378)
- SqlPrimaryExpression (page 376)
- WhereClause (page 394)
- WindowFunctionCall (page 390)

9.3.2.2.19 SqlAndExpression




Definition:

SqlAndExpression ::= SqlUnaryLogicalExpression ('AND' SqlUnaryLogicalExpression)*

Used In:

- SqlExpression (page 376)

9.3.2.2.20 SqlUnaryLogicalExpression




Definition:

SqlUnaryLogicalExpression ::= 'NOT'? SqlRelationalExpression

Used In:

- SqlAndExpression (page 377)

9.3.2.2.21 SqlRelationalExpression



Definition:

```
SqlRelationalExpression ::= ( '(' SqlExpressionList ')' | SqlSimpleExpression )
( SqlRelationalOperatorExpression | SqlInClause | SqlBetweenClause |
SqlLikeClause | IsNullClause )?
```

Used In:

- `SqlUnaryLogicalExpression` (page 377)

9.3.2.2.22 SqlExpressionList




Definition:

```
SqlExpressionList ::= SqlExpression ( ',' SqlExpression )*
```

Used In:

- `ArrayValueConstructor` (page 391)
- `GenericFunctionCall` (page 387)
- `GroupingOperation` (page 390)
- `SqlInClause` (page 379)
- `SqlRelationalExpression` (page 378)
- `WindowSpecification` (page 388)

9.3.2.2.23 SqlRelationalOperatorExpression




Definition:

SqlRelationalOperatorExpression ::= Relop ('(' QueryExpressionBody ')' | SqlSimpleExpression)

Used In:

- SqlRelationalExpression (page 378)

9.3.2.2.24 Relop




Definition:

Relop ::= '=' | '!=' | '#' | '<>' | '>' | '>=' | '<' | '<='

Used In:

- SqlRelationalOperatorExpression (page 378)

9.3.2.2.25 SqlInClause



Definition:

SqlInClause ::= 'NOT'? 'IN' '(' SqlExpressionList ')'

Used In:

- SqlRelationalExpression (page 378)

9.3.2.2.26 SqlBetweenClause




Definition:

SqlBetweenClause ::= 'NOT'? 'BETWEEN' SqlSimpleExpression 'AND' SqlSimpleExpression

Used In:

- `SqlRelationalExpression` (page 378)

9.3.2.2.27 SqlLikeClause




Definition:

SqlLikeClause ::= 'NOT'? 'LIKE' SqlSimpleExpression

Used In:

- `SqlRelationalExpression` (page 378)

9.3.2.2.28 IsNullClause




Definition:

IsNullClause ::= 'IS' 'NOT'? 'NULL'

Used In:

- `SqlRelationalExpression` (page 378)

9.3.2.2.29 Constant




Definition:

Constant ::= 'NULL' | ('+' | '-')? S_NUMBER | S_CHAR_LITERAL | 'TRUE' | 'FALSE' | S_PARAMETER

Used In:

- ConstantList (page 387)
- SqlPrimaryExpression (page 376)
- WindowFunctionCall (page 390)

9.3.2.2.30 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 376)

9.3.2.2.31 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 382)

9.3.2.2.32 UserDefinedType




Definition:

UserDefinedType ::= S_IDENTIFIER ('.' S_IDENTIFIER)?

Used In:

- CastableDataTypeDeclaration (page 384)

9.3.2.2.33 SimpleCaseCall



Definition:

SimpleCaseCall ::= SqlSimpleExpression ('WHEN' SqlSimpleExpression 'THEN' SqlSimpleExpression)+ ('ELSE' SqlSimpleExpression)? 'END'

Used In:

- FunctionCall (page 382)

9.3.2.2.34 BooleanCaseCall




Definition:

BooleanCaseCall ::= ('WHEN' SqlExpression 'THEN' SqlSimpleExpression)+ ('ELSE' SqlSimpleExpression)? 'END'

Used In:

- FunctionCall (page 382)

9.3.2.2.35 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 376)

9.3.2.2.36 SourceField



Definition:

SourceField ::= ObjectName ('.' ObjectName ('.' ObjectName)?)?

Used In:

- AggregateFunctionCall (page 386)
- SourceFieldList (page 394)
- WindowFunctionCall (page 390)

9.3.2.2.37 GenericFunctionCall




Definition:

GenericFunctionCall ::= ObjectName '(' SqlExpressionList? (';' ConstantList)? ')'

Used In:

- SqlPrimaryExpression (page 376)

9.3.2.2.38 ConstantList




Definition:

ConstantList ::= Constant (',' Constant)*

Used In:

- GenericFunctionCall (page 387)

9.3.2.2.39 WindowSpecification




Definition:

WindowSpecification ::= 'OVER' '(' ('PARTITION' 'BY' SqlExpressionList)? OrderByClause? WindowFrameClause? ')'

Used In:

- SqlPrimaryExpression (page 376)
- WindowFunctionCall (page 390)

9.3.2.2.40 OrderByClause




Definition:

OrderByClause ::= 'ORDER' 'BY' SqlExpression ('ASC' | 'DESC')? ('NULLS' ('FIRST' | 'LAST'))? (',' SqlExpression ('ASC' | 'DESC')? ('NULLS' ('FIRST' | 'LAST'))?)*

Used In:

- SelectWithOrder (page 371)
- WindowSpecification (page 388)

9.3.2.2.41 WindowFrameClause




Definition:

WindowFrameClause ::= ('ROWS' | 'RANGE') (WindowFrameStart | 'BETWEEN' WindowFrameBound 'AND' WindowFrameBound)

Used In:

- WindowSpecification (page 388)

9.3.2.2.42 WindowFrameStart




Definition:

WindowFrameStart ::= ('UNBOUNDED' | FrameBound) 'PRECEDING' | 'CURRENT' 'ROW'

Used In:

- WindowFrameBound (page 389)
- WindowFrameClause (page 388)

9.3.2.2.43 FrameBound




Definition:

FrameBound ::= S_NUMBER | S_PARAMETER

Used In:

- WindowFrameBound (page 389)
- WindowFrameStart (page 389)

9.3.2.2.44 WindowFrameBound




Definition:

WindowFrameBound ::= ('UNBOUNDED' | FrameBound) 'FOLLOWING' | WindowFrameStart

Used In:

- WindowFrameClause (page 388)

9.3.2.2.45 GroupingOperation




Definition:

GroupingOperation ::= 'GROUPING' '(' SqlExpressionList ')'

Used In:

- SqlPrimaryExpression (page 376)

9.3.2.2.46 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 376)

9.3.2.2.47 LongConstant




Definition:

LongConstant ::= S_NUMBER | S_PARAMETER

Used In:

- LimitClause (page 398)
- WindowFunctionCall (page 390)

9.3.2.2.48 IntegerConstant




Definition:

IntegerConstant ::= S_NUMBER | S_PARAMETER

Used In:

- ArrayElementReference (page 391)
- WindowFunctionCall (page 390)

9.3.2.2.49 ArrayElementReference



Definition:

ArrayElementReference ::= '[' IntegerConstant ']'

Used In:

- SqlPrimaryExpression (page 376)

9.3.2.2.50 ArrayValueConstructor




Definition:

ArrayValueConstructor ::= 'ARRAY' '[' SqlExpressionList ']'

Used In:

- SqlPrimaryExpression (page 376)

9.3.2.2.51 Alias




Definition:

Alias ::= 'AS'? ObjectName

Used In:

- SelectItem (page 373)
- SourceReferenceOrSubquery (page 392)

9.3.2.2.52 FromClause




Definition:

FromClause ::= 'FROM' FromItem (',' FromItem)*

Used In:

- SelectWithoutOrder (page 372)

9.3.2.2.53 FromItem




Definition:

FromItem ::= SourceReferenceOrSubquery JoinedTable*

Used In:

- FromClause (page 392)

9.3.2.2.54 SourceReferenceOrSubquery




Definition:

SourceReferenceOrSubquery ::= (SourceReference | '(' SelectWithOrder ')') Alias?

Used In:

- FromItem (page 392)
- JoinedTable (page 393)

9.3.2.2.55 SourceReference




Definition:

SourceReference ::= ObjectName ('.' ObjectName)? ('/' ObjectName ('.' ObjectName)?)*

Used In:

- SourceReferenceOrSubquery (page 392)

9.3.2.2.56 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 392)

9.3.2.2.57 SourceFieldList



Definition:

SourceFieldList ::= SourceField (',' SourceField)*

Used In:

- JoinedTable (page 393)

9.3.2.2.58 WhereClause




Definition:

WhereClause ::= 'WHERE' SqlExpression

Used In:

- SelectWithoutOrder (page 372)

9.3.2.2.59 GroupByClause




Definition:

GroupByClause ::= 'GROUP' 'BY' ('ALL' | 'DISTINCT')? GroupingElementList

Used In:

- SelectWithoutOrder (page 372)

9.3.2.2.60 GroupingElementList




Definition:

GroupingElementList ::= GroupingElement (',' GroupingElement)*

Used In:

- GroupByClause (page 394)

9.3.2.2.61 GroupingElement




Definition:

GroupingElement ::= EmptyGroupingSet | OrdinaryGroupingSet | RollupList | CubeList | GroupingSetsSpecification

Used In:

- GroupingElementList (page 394)

9.3.2.2.62 EmptyGroupingSet




Definition:

EmptyGroupingSet ::= '(' ')'

Used In:

- GroupingElement (page 395)
- GroupingSet (page 398)

9.3.2.2.63 OrdinaryGroupingSet



Definition:

OrdinaryGroupingSet ::= '(' GroupingExpressionReferenceList ')' | GroupingExpressionReference

Used In:


Definition:

OrdinaryGroupingSetList ::= OrdinaryGroupingSet (',' OrdinaryGroupingSet)*

Used In:

- CubeList (page 397)
- RollupList (page 396)

9.3.2.2.68 CubeList




Definition:

CubeList ::= 'CUBE' '(' OrdinaryGroupingSetList ')'

Used In:

- GroupingElement (page 395)
- GroupingSet (page 398)

9.3.2.2.69 GroupingSetsSpecification




Definition:

GroupingSetsSpecification ::= 'GROUPING' 'SETS' '(' GroupingSetList ')'

Used In:

- GroupingElement (page 395)
- GroupingSet (page 398)

9.3.2.2.70 GroupingSetList




Definition:

GroupingSetList ::= GroupingSet (',' GroupingSet)*

Used In:

- GroupingSetsSpecification (page 397)

9.3.2.2.71 GroupingSet



Definition:

GroupingSet ::= EmptyGroupingSet | OrdinaryGroupingSet | RollupList | CubeList | GroupingSetsSpecification

Used In:

- GroupingSetList (page 397)

9.3.2.2.72 HavingClause




Definition:

HavingClause ::= 'HAVING' SqlExpression

Used In:

- SelectWithoutOrder (page 372)

9.3.2.2.73 LimitClause




Definition:

LimitClause ::= 'LIMIT' LongConstant (',' LongConstant)?

Used In:

- SelectWithOrder (page 372)

9.3.2.2.74 S_NUMBER




Definition:

S_NUMBER ::= FLOAT ([eE] [-+]? INTEGER)?

Used In:

- Constant (page 381)
- FrameBound (page 389)
- IntegerConstant (page 391)
- LongConstant (page 391)

9.3.2.2.75 FLOAT




Definition:

Float ::= (INTEGER? '.')? INTEGER

Used In:

- S_NUMBER (page 399)

9.3.2.2.76 INTEGER



Definition:

INTEGER ::= DIGIT+

Used In:

- FLOAT (page 399)
- S_NUMBER (page 399)

9.3.2.2.77 DIGIT




Definition:

DIGIT ::= [0-9]

Used In:

- INTEGER (page 399)

9.3.2.2.78 S_PARAMETER




Definition:

S_PARAMETER ::= ':' (S_IDENTIFIER | S_QUOTED_IDENTIFIER)

Used In:

- CharacterConstant (page 370)
- Constant (page 381)
- FrameBound (page 389)
- IntegerConstant (page 391)
- LongConstant (page 391)

9.3.2.2.79 S_IDENTIFIER




Definition:

S_IDENTIFIER ::= NAMESTARTCHAR NAMECHAR*

Used In:

- ObjectName (page 373)
- S_PARAMETER (page 400)
- UserDefinedType (page 385)

9.3.2.2.80 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)

9.3.2.2.81 NAMECHAR



Definition:

NAMECHAR ::= NAMESTARTCHAR | [:\$0-9#x00B7#x0300-#x036F#x203F-#x2040]

Used In:

- S_IDENTIFIER (page 400)

9.3.2.2.82 S_QUOTED_IDENTIFIER




Definition:

S_QUOTED_IDENTIFIER ::= '"' [^"#\x000A#\x000D]* '''

Used In:

- ObjectName (page 373)
- S_PARAMETER (page 400)

9.3.2.2.83 S_CHAR_LITERAL




Definition:

S_CHAR_LITERAL ::= ''' (''' '''' | [^']) * '''

Used In:

- CharacterConstant (page 370)
- Constant (page 381)


9.3.2.2.84 IgnorableWhitespace



Definition:

IgnorableWhitespace ::= [#x0009#x000D#x000A] | LINE_COMMENT | MULTI_LINE_COMMENT

9.3.2.2.85 LINE_COMMENT




Definition:

LINE_COMMENT ::= '--' [^#x000D#x000A] *

Used In:

- IgnorableWhitespace (page 403)

9.3.2.2.86 MULTI_LINE_COMMENT



Definition:

MULTI_LINE_COMMENT ::= '/' '*' ([^*] | '*' + [^*/]) * '*' + '/'

Used In:

- IgnorableWhitespace (page 403)

9.3.3 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 420) for instructions.

Note that in contrast to the case-insensitive handling of SQL functions, RAQL-specific functions are handled case-sensitively.

9.3.3.1 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 382) 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 417) for a complete list and the alternate functions from version 3.8 to use as replacements.

9.3.3.1.1 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 ▪ day-of-week ▪ hour ▪ minute

Function	Description
	<ul style="list-style-type: none"> ▪ 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(Date value) format_date(Date value, String datePattern) format_date(Date value, String datePattern, String timeZone)</p>	<p>Converts a date value to a string representation in given target format and target timezone.</p> <ul style="list-style-type: none"> ▪ datePattern may be 'short', 'medium', 'long', 'full' or any date pattern that is accepted by Java's SimpleDateFormat. The variant without datePattern parameter uses this pattern: yyyy-MM-dd"T"HH:mm:ss.SSSZ

Function	Description
	<ul style="list-style-type: none"> ▪ timeZone may be any time zone ID that is understood by Java's TimeZone.getTimeZone method. For more details, see the corresponding Java documentation. Variants of this function without a timeZone parameter use the configured feed processing time zone (page 153). <p>Examples</p> <ul style="list-style-type: none"> ▪ <code>format_date(value)</code> ▪ <code>format_date(value, 'yyyy-MM-dd')</code> ▪ <code>format_date(value, 'h:mm a', 'GMT-08:00')</code> ▪ <code>format_date(value, 'HH:mm:ss', 'Europe/Berlin')</code>
<code>format_date (dateColumn, datePattern)</code>	Corresponds to format_date(dateColumn, datePattern, 'GMT')
<code>format_date (dateColumn)</code>	Corresponds to format_date(dateColumn, 'yyyy-MM-dd'T'HH:mm:ss.SSSZ', 'GMT')
<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>	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). For example: where matches(firstname,'ra[n l])='true' Will match Frank, Gerald or Randy in the firstname column. To perform exact regular expression matches, with no implicit wildcards, see the regex function.
<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>	Replaces null values in the specified column or expression with the specified string value. 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.

Function	Description
project(Xcol_1, X col_2, ..., X col_n, Number index)	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.
quarter(Date col-or-expr) quarter(String col-or-expr) quarter(Long col-or-expr)	Converts the values for this column or expression to a date and returns the number of the quarter as an integer.
regex(String col-or-expr, String pattern)	Determines if a string matching the regular expression in pattern exists in the values of the column or expression as specified. Matching is case insensitive. Returns a boolean value (true or false). 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. For example: where regex(firstname,'ra[n l]')=true Will match Ralph or Randy in the firstname column, but will not match Frank or Gerald.
split_part(String col-or-expr, String delim, Integer part) split_part(String col-or-expr, Character delim, Integer part)	Returns the n th part of the values of this column or expression when each value is split into parts at each delimiter defined in delim . For example: split_part("INV:2012:GHI345",":",3) Returns the string "GHI345".
split_regex((String col-or-expr, String pattern, Integer part)	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.
time_mask(Date column, String mask) time_mask(String column, String mask)	Converts the value for this column or expression to a date and time, if needed, in the form yyyy-MM-dd:HH:mm:ss . Note: If the values for the column do not contain times, the time is set to 00:00:00 . It extracts the portion of this date and time through the time part identified by the mask. Valid masks include: <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.

Function	Description
	<ul style="list-style-type: none"> ▪ m = returns the full date, hour and minute. ▪ s = returns the full date and time.
<p>to_date(String value)</p> <p>to_date(String value, String datePattern)</p> <p>to_date(String value, String datePattern, String timeZone)</p>	<p>Converts a string value to Date.</p> <p>The datePattern, if present, is any pattern that is accepted by Java's SimpleDateFormat. If no pattern is specified, this function can have a negative effect on performance because it attempts to convert the string iterating over the following date formats:</p> <ul style="list-style-type: none"> ▪ yyyy-MM-dd HH:mm:ss.SSS ▪ yyyy-MM-dd HH:mm:ss ▪ yyyy-MM-dd'T'HH:mm:ss.SSSXXX ▪ yyyy-MM-dd'T'HH:mm:ss.SSS ▪ yyyy-MM-dd'T'HH:mm:ssXXX ▪ yyyy-MM-dd'T'HH:mm:ss ▪ M/dd/yy or the short format for the current locale ▪ MMMM d, yyyy or the long format for the current locale ▪ EEEE, MMMM d, yyyy or the full format for the current locale ▪ EEE MMM d HH:mm:ss z yyyy ▪ EEE, d MMM yyyy HH:mm:ss Z ▪ dd.MM.yyyy HH:mm:ss ▪ yyyy-MM-dd ▪ dd-MM-yyyy ▪ dd.MM.yy ▪ MM/dd/yy <p>Use the datePattern 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> <p>The timeZone may be any time zone ID that is understood by Java's TimeZone.getTimeZone method. For more details, see the corresponding Java documentation. Variants of this function without a timeZone parameter use the configured feed processing time zone (page 153).</p>

Function	Description
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.

9.3.3.1.2 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 386) 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 419) 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"> ▪ $< =$ minimum value for the column, it returns 0. ▪ $> =$ 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)	<p>Returns the geometric mean of the values for this column or expression within the partition or the current window in this 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>
✓		analytics.kmeans_clusters(Number col-or-expr1[,Number col-or-expr2,...Number col-or-	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

Aggregate	Window	Function	Description
		exprN]; Integer k, Integer iterations, String measure)	<p>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 ▪ squaredeuclidean
	✓	analytics.kmeans_observations(Number col-or-expr1[,Number col-or-expr2,...Number col-or-exprN]; Integer k, Integer iterations, String measure)	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

Aggregate	Window	Function	Description
			<p>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: ▪ 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 ▪ squaredeuclidean
	✓	analytics.ordinal(String col-or-expr)	Returns an index number for each unique 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)	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

Aggregate	Window	Function	Description
			<p>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>

9.3.3.1.3 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]

rowNumber	ROWS BETWEEN 3 PRECEDING AND 3 FOLLOWING	ROWS BETWEEN 3 PRECEDING AND 1 PRECEDING	ROWS BETWEEN CURRENT ROW AND UNBOUNDED FOLLOWING
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]	
5	2400	2300	[2100,2700]	
6	2750	2866,67	[2450,3050]	
7	2900	2980	[2600,3200]	

rowNumber	salary	SELECT avg(salary) OVER (ORDER BY salary RANGE BETWEEN 300 PRECEDING AND 300 FOLLOWING) FROM ...	considered window	
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.

9.3.3.1.4 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.

9.3.3.1.4.1 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>
"second"(Date column) "second"(String column) "second"(Long column)	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)
"week"(Date column) "week"(String column) "week"(Long column)	extract_week
"year"(Date column) "year"(String column) "year"(Long column) "year"(Object column)	extract_year

9.3.3.1.4.2 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)
covariance(Number column, Number column)	covar(Number column, Number column)

Deprecated or Unsupported 3.7 Functions	Replace with 3.8 Function
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()

9.3.3.2 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 421).

Then:

1. Write Plain Functions for RAQL (page 422)
2. Write Window Analytic Functions for RAQL (page 427)
3. Write Aggregate Analytic Functions for RAQL (page 433)
4. Use a Factory for Function Overloading (page 439), if needed
5. Configure, Compile, Deploy and Test User-Defined Functions (page 424)

9.3.3.2.1 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 424).

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 424) 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 424).

9.3.3.2.2 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 426) 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.

9.3.3.2.3 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 425) 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-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 121) (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.

9.3.3.2.3.1 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" } ] }
    ]
  }
}
```

9.3.3.2.3.2 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 121) for instructions.

9.3.3.2.4 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 421) 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 427)
2. [Implement the call Method for the Window Analytic Function](#) (page 428)

And [Configure, Compile, Deploy and Test User-Defined Functions](#) (page 424). For the complete code, see [Complete WindowSum Example](#) (page 430).

The second example implements the MashZone NextGen built-in lead function which illustrates techniques to [Work with Specific Records in Window Calculations](#) (page 432) using the current position of a record.

9.3.3.2.4.1 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.

9.3.3.2.4.2 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 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
}
```

```
}  
}
```

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.

9.3.3.2.4.3 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
}
}

```

9.3.3.2.4.4 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.

9.3.3.2.5 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 421) for details. Write the class for your aggregate analytic function and then Configure, Compile, Deploy and Test User-Defined Functions (page 424).

We're going to use two examples. The first example, My Average Aggregate Example (page 434) 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 436), 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.

9.3.3.2.5.1 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 == 0l)
        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.

9.3.3.2.5.2 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.

9.3.3.2.6 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 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 investigate the symptom.

9.4.1 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 inspect your global sales processes of the last year and you find that a disproportionately large number of customer complaints occur in your subsidiary in **China** compared to 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
- 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.

9.4.2 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 up to 50 results, consisting of the top 25 causes that promote the symptom and the top 25 causes that inhibit the symptom.

9.4.3 Score

The score output by PPM is a measure of the association between a process with the symptom and also with 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 439) 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.

9.4.4 Scaled system

The **Root Cause Miner** feature is also available in a scaled system. The sub-servers 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 sub-servers and much network traffic between the master and its sub-servers.

9.4.5 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.

9.4.6 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 106).
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 using the consulting services provided by Software AG, 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 describe only 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 Support

If you have any questions on specific installations that you cannot perform yourself, contact your local Software AG sales organization (<https://www.softwareag.com/corporate/company/global/offices/default.html>). To get detailed information and support, use our websites.

If you have a valid support contract, you can contact **Global Support ARIS** at: **+800 ARISHelp**. If this number is not supported by your telephone provider, please refer to our Global Support Contact Directory.

ARIS COMMUNITY

Find information, expert articles, issue resolution, videos, and communication with other ARIS users. If you do not yet have an account, register at ARIS Community.

SOFTWARE AG EMPOWER PORTAL

You can find documentation on the Software AG Documentation website (<https://empower.softwareag.com/>). The site requires credentials for Software AG's Product

Support site **Empower**. If you do not yet have an account for **Empower**, send an e-mail to empower@softwareag.com (mailto:empower@softwareag.com) with your name, company, and company e-mail address and request an account.

If you have no account, you can use numerous links on the TECHcommunity website. For any questions, you can find a local or toll-free number for your country in our Global Support Contact Directory and give us a call.

TECHCOMMUNITY

On the **TECHcommunity** website, you can find documentation and other technical information:

- Use the online discussion forums, moderated by Software AG professionals, to ask questions, discuss best practices, and learn how other customers are using Software AG technology.
- Access articles, code samples, demos, and tutorials.
- Find links to external websites that discuss open standards and web technology.
- Access product documentation, if you have **TECHcommunity** credentials. If you do not, you will need to register and specify **Documentation** as an area of interest.

EMPOWER (LOGIN REQUIRED)

If you have an account for **Empower**, use the following sites to find detailed information or get support:

- You can find product information on the Software AG Empower Product Support website.
- To get information about fixes and to read early warnings, technical papers, and knowledge base articles, go to the Knowledge Center.
- Once you have an account, you can open Support Incidents online via the eService section of Empower.
- To submit feature/enhancement requests, get information about product availability, and download products, go to Products.

SOFTWARE AG MANAGED LEARNINGS

Get more information and trainings to learn from your laptop computer, tablet or smartphone. Get the knowledge you need to succeed and make each and every project a success with expert training from Software AG.

If you do not have an account, register as a customer or as a partner.